

HP-UX System Administrator's Guide: Configuration Management

HP-UX 11i Version 3

HP Part Number: 5992-6579
Published: September 2009
Edition: 4



© Copyright 2008, 2009 Hewlett-Packard Development Company L.P.

Legal Notices

Proprietary computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Warranty The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett-Packard shall not be held liable for technical or editorial errors or omissions contained herein.

A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

Trademark Notices The following trademarks are acknowledged:

Intel® and Itanium® are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

MS-DOS®, Windows NT®, and Microsoft® are registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

X Window System™ is a trademark of X Consortium, Inc.

Table of Contents

Preface.....	11
Intended Audience.....	11
About This Series.....	11
About This Document.....	11
Finding HP-UX Information.....	12
HP-UX 11i Release Names and Operating System Version Identifiers.....	13
Determining Your System Version.....	14
Typographic Conventions.....	15
Examples and Shells.....	16
Command Syntax.....	16
Function Synopsis and Syntax.....	16
Publication History.....	17
 1 Introduction.....	19
Configuration Tools.....	19
HP Systems Insight Manager.....	19
HP-UX System Management Homepage (Web-Based).....	20
HP-UX System Management Homepage (Text-Based).....	20
HP-UX Commands.....	20
Using HP SIM and HP SMH versus HP-UX Commands.....	20
HP-UX System Management Homepage (Web-Based).....	21
Starting an HP SMH Web Client.....	24
Starting the HP SMH Daemon on the Target System.....	25
HP-UX System Management Homepage (Text-Based).....	25
Starting Text-Based HP SMH.....	30
Giving Users Limited Access to Text-Based HP SMH.....	31
Using the X Window System.....	32
Setup for an HP-UX Local System.....	32
Setup for a Non-HP-UX Client System.....	33
 2 Configuring System Parameters.....	35
Starting A Preloaded System.....	35
Preventing Users from Logging In.....	36
Changing Login Special Characters.....	37
Controlling Usage and Processes with Run-Levels.....	38
Setting the System Clock.....	40
Potential Problems When Changing the System Clock.....	40
Setting the Time Zone (TZ).....	40
Setting the Time and Date.....	40

Setting System and Network Parameters.....	41
Customizing System-Wide and User Login Environments.....	42
Setting Shadow Password Mode.....	43
To Switch to Shadow Password Mode.....	43
To Switch to Standard Password Mode.....	43
Setting Long User, Group, Host, and File Names.....	43
Setting Long User and Group Names.....	43
Setting Long Host Names.....	44
Setting Long File Names.....	44
Configuring /etc/hosts.....	44
Configuring New HP-UX Systems into Workgroups.....	45
Configuring a New System into a Network.....	45
 3 Configuring Root.....	47
Create the Root Home Directory.....	47
 4 Configuring Users and Groups.....	49
Configuring Users and Groups with HP SMH.....	49
Adding a User with Text-Based HP SMH.....	50
Skeleton Directory.....	54
Changing the Skeleton Directory.....	54
Using a Template to Add a User with Text-Based HP SMH.....	55
Making User Templates with Text-Based HP SMH.....	55
Modifying a User with Text-Based HP SMH.....	57
Deleting a User with Text-Based HP SMH.....	58
Adding a Group with Text-Based HP SMH.....	59
Managing Users and Groups with Commands.....	60
Automating the Process of Adding a User.....	60
Using the useradd, usermod, and userdel Commands.....	60
Manually Adding a User.....	61
Configuring System and User Security.....	63
Configuring System Default Security Attributes.....	63
Configuring User Security Attributes.....	65
Controlling File Access.....	67
Defining Group Membership.....	67
Setting File Access Permissions.....	68
Setting Ownership for Files.....	68
Setting Access Control Lists.....	68
Customizing System-Wide and User Login Environments.....	68
Accessing Multiple Systems.....	69
Sharing Remote Work Directories.....	69
Local versus Remote Home Directories.....	70
Adding a User to Several Systems: A Case Study.....	70

Sharing a Local Home Directory.....	72
Allowing Access from Remote Systems.....	72
\$HOME/.rhosts File.....	72
Enabling X Window Server Access.....	73
 5 Configuring Networking.....	75
Configuring the Network File System (NFS).....	75
Exporting versus Sharing.....	75
Enabling and Disabling the NFS Server and Client.....	76
Enable or Disable the NFS Server with Text-Based HP SMH.....	76
Enable or Disable the NFS Client with Text-Based HP SMH.....	76
Sharing an HP-UX Directory.....	77
Share a Directory Using Text-Based HP SMH.....	77
Share a Directory Using HP-UX Commands.....	77
Mounting a Shared File System (HP-UX to HP-UX).....	78
Standard-Mount a Shared File System Using Text-Based HP SMH.....	80
Automount a Shared File System Using Text-Based HP SMH.....	82
Standard-Mount a Shared File System Using HP-UX Commands.....	82
Troubleshooting NFS.....	84
Requisite Entries.....	85
Recovering Network Services after a Power Failure.....	86
Symptoms and Keywords.....	86
What To Do.....	86
Moving or Reusing a Shared Directory.....	87
Configuring File Transfer Protocol (FTP).....	87
Configuring Anonymous FTP.....	87
Configuring Anonymous FTP with Text-Based HP SMH.....	87
Configuring Anonymous FTP with HP-UX Commands.....	88
Troubleshooting FTP Login.....	90
Setting Up /etc/shells.....	90
Configuring HP-UX Systems for File Transfer.....	91
FTP (File Transfer Protocol).....	91
Establishing an FTP Connection from HP-UX to a PC.....	91
On the HP-UX System: Retrieving a File from the PC.....	95
On the HP-UX System: Sending a File to the PC.....	96
Establishing an FTP Connection from a PC to HP-UX.....	99
On the PC: Retrieving a File from the HP-UX System.....	102
On the PC: Sending a File to the HP-UX System.....	104
Interfacing with Microsoft Windows Systems.....	106
Hardware Connections.....	106
Sharing Directories between HP-UX and MS Windows.....	106
HP CIFS.....	107
Configuring HP-UX Systems for Terminal Emulation.....	107

Telnet.....	107
Using Telnet to Log in to a PC from an HP-UX System.....	108
Using Telnet to Log in to an HP-UX System from a PC.....	109
Other Terminal Emulators.....	110
 6 Configuring Printers.....	111
Starting and Stopping the LP Spooler.....	111
Starting and Stopping Spooler Using Text-Based HP SMH.....	112
Starting and Stopping Spooler Using HP-UX Commands.....	112
Adding a Local Printer to the LP Spooler.....	112
Adding a Local Printer Using HP SMH.....	112
Adding a Local Printer Using HP-UX Commands.....	112
Adding a Remote Printer to the LP Spooler.....	113
Adding a Remote Printer Using Text-Based HP SMH.....	113
Adding a Remote Printer Using HP-UX Commands.....	114
Adding a Network Printer with HP JetDirect.....	116
Adding A Network Printer Using Web-Based HP SMH.....	116
Adding A Network Printer Using HP-UX Commands.....	116
Creating a Printer Class.....	117
Creating a Printer Class Using HP-UX Commands.....	117
Removing a Printer from the LP Spooler.....	118
Removing a Printer Using Web-Based HP SMH.....	118
Removing a Printer Using HP-UX Commands.....	118
Removing a Printer from a Printer Class.....	119
Removing a Printer from a Class Using HP-UX Commands.....	119
Removing a Printer Class.....	120
Removing a Printer Class Using HP-UX Commands.....	120
 7 Configuring Mail.....	121
Components of an Electronic Mail System.....	121
Mail User Agents.....	121
Mail Delivery Agents.....	121
Mail Alias Files.....	122
The Mail Queue.....	122
Networking Topographies.....	123
Central Mail Hub.....	123
Gateway Mail Hub.....	123
Fully Distributed.....	124
Selecting a Topography.....	124
MIME Applications.....	125
Configuring a System to Send Electronic Mail.....	125
Using Site Hiding.....	125
Configuring a System to Receive Electronic Mail.....	126

Central Mail Hub Topography (Receiving E-mail).....	126
Configuring the Hub.....	127
Configuring the Clients.....	127
Gateway Mail Hub Topography (Receiving E-mail).....	127
Configuring the Hub.....	128
Configuring the Clients.....	128
Fully Distributed (Standalone System) Topography.....	130
Configuring Each System.....	130
 8 Setting Up the Online Manpages.....	131
 9 Configuring Peripherals.....	133
Configuring PCI Error Recovery.....	133
What is PCI Error Recovery?.....	134
Controlling PCI Error Recovery.....	134
PCI Error Recovery Documentation.....	135
Ethernet Configuration and Verification.....	135
Installing or Updating a Networking Driver with swinstall.....	136
Assigning IP Addresses and Subnet Masks to Additional Cards Using Web-Based HP SMH.....	137
Frequently Asked Questions.....	137
Verifying LAN Installation.....	138
Routing: Configuring the LAN to Reach Other Networks.....	139
Setting Speed and Duplex Mode of the LAN Card and Link Partner (Switch or Router).....	140
Ensuring Card and Link Partner Speed and Duplex Settings.....	140
Configuring Optional Jumbo Frames Size for Gigabit Ethernet.....	141
Verifying that the Networking Driver is in the Kernel.....	142
Related System and Network Administration.....	142
Further Ethernet Information.....	143
Configuring Non-HP Terminals.....	143
Troubleshooting Terminals.....	145
Unresponsive Terminals.....	145
Other Terminal Problems.....	151
What to Check for When Garbage is Mixed with Valid Data.....	151
What to Check for When Everything Printed is Garbage.....	151
 10 Configuring the Kernel.....	153
Kernel Configuration Features.....	153
What Is a Kernel Configuration?.....	154
Overview of Kernel Configuration Commands.....	154
Overview of HP SMH for Kernel Configuration.....	155

Other Kernel Configuration Operations.....	156
Common Behavior for Kernel Configuration Commands.....	157
Common Command Line Options.....	157
Common Output Formats.....	158
Common Exit Status Codes.....	159
Common Security Constraints.....	159
Persistence of Changes.....	159
Managing Kernel Modules with kcmodule.....	160
Getting Information About Modules.....	160
With No Options.....	160
With the -d Option.....	161
With the -v Option.....	161
With the -P Option.....	161
Interpreting Module Information.....	162
Changing Module States.....	164
Managing Kernel Modules with HP SMH.....	165
Getting Information about Modules.....	166
Interpreting Module Information.....	167
Changing Module States.....	168
Managing Kernel Tunable Parameters with kctune.....	170
Getting Information About Tunables.....	171
Interpreting Tunable Information.....	173
Changing Tunable Values.....	174
Managing Kernel Tunable Parameters with HP SMH.....	175
Getting Information About Tunables.....	176
Interpreting Tunable Information.....	177
Changing Tunable Values.....	178
Monitoring Kernel Resource Usage.....	180
Getting Information about Alarms.....	181
Interpreting Alarm Information.....	182
Changing Alarm Settings.....	183
Add an Alarm.....	186
Activate, Deactivate, or Remove an Alarm.....	187
Resource Usage Commands.....	188
Managing the Running Configuration Using kconfig.....	188
Managing Saved Configurations with kconfig.....	189
Getting Information about Saved Configurations.....	189
Interpreting Saved Configuration Information.....	190
Using and Modifying Saved Configurations.....	191
Creating Saved Configurations.....	191
Using Saved Configurations.....	191
Modifying Saved Configurations.....	192
Managing Configurations with System Files.....	192
Making Configuration Changes with System Files.....	192

Uses for System Files.....	194
The Kernel Configuration Log File.....	195
Parsing Command Output.....	196
Recovering from Errors.....	197
The Automatic Backup Configuration.....	198
Booting a Saved Configuration.....	198
Booting in Fail-Safe Mode.....	199
Modifying Tunable Values at Boot Time.....	200
Guidelines for Recovering from Errors.....	201
Kernel Configuration Example.....	201
Kernel Configuration Quick Reference Tables.....	207
Transition from Previous HP-UX Releases.....	210
 Index.....	215

Preface

Intended Audience

The *HP-UX System Administrator's Guide* series is written for administrators of HP-UX systems of all skill levels needing to administer HP-UX systems, beginning with HP-UX Release 11i Version 3.

While many topics in this set apply to previous releases, much has changed in HP-UX 11i Version 3. Therefore, for information about prior releases, please refer to *Managing Systems and Workgroups: A Guide for System Administrators*.

About This Series

The *HP-UX System Administrator's Guide* documents the core set of tasks (and associated concepts) necessary to administer HP-UX 11i Version 3.

The *HP-UX System Administrator's Guide* is a set of documents, comprised of the following volumes:

<i>Overview</i>	Provides a high-level view of HP-UX 11i Version 3, its components, and how they relate to each other.
<i>Configuration Management</i>	Describes many of the tasks you need to perform to configure and customize system settings and the behavior of subsystems.
<i>Logical Volume Management</i>	Documents how to configure physical volumes, volume groups, and logical volumes using the HP Logical Volume Manager (LVM).
<i>Security Management</i>	Documents the data and system security features of HP-UX 11i Version 3.
<i>Routine Management Tasks</i>	Documents many of the ongoing tasks you need to perform to keep your system running smoothly.

About This Document

HP-UX System Administrator's Guide: Configuration Management describes the tools and processes for configuring a system and connecting it to the network. It includes the following major topics:

<i>Chapter 1 "Introduction"</i>	General configuration topics including an overview of the major configuration tools, HP Systems Insight Manager (HP SIM), and HP System Management Homepage (HP SMH).
---------------------------------	---

Chapter 2: "Configuring System Parameters"	Starting a preinstalled system and general initialization.
Chapter 3: "Configuring Root"	Setting up a root user home directory and configuring root user parameters.
Chapter 4: "Configuring Users and Groups"	Adding users to the system.
Chapter 5: "Configuring Networking"	Setting up the network interface, sharing directories with other systems, transferring files between systems.
Chapter 6: "Configuring Printers"	Configuring local, remote, and network printers.
Chapter 7: "Configuring Mail"	Setting up various mail server configurations.
Chapter 8: "Setting Up the Online Manpages"	Managing disk usage versus access time.
Chapter 9: "Configuring Peripherals"	PCI recovery, Ethernet configuration, terminal troubleshooting,
Chapter 10: "Configuring the Kernel"	Managing kernel tunables, modules, and alarms.

For other configuration topics, see:

- *System Administrator's Guide: Security Management*
- *System Administrator's Guide: Logical Volume Management*

Finding HP-UX Information

Table 1 outlines where to find general system administration information for HP-UX. This table does not include information for specific products.

Table 1 Finding HP-UX Information

If you need to . . .	Refer to . . .	Located at . . .
Find out: <ul style="list-style-type: none">• What has changed in HP-UX releases• The content of the Operating Environments• Firmware requirements, and supported systems for a specific release	The HP-UX 11i Release Notes specific to your version of HP-UX.	<ul style="list-style-type: none">• HP Instant Information• HP Technical Documentation web site http://docs.hp.com
Install or update HP-UX	<ul style="list-style-type: none">• <i>Read Before Installing or Updating to HP-UX</i>• <i>HP-UX 11i Installation and Update Guide</i>	<ul style="list-style-type: none">• Media Kit (supplied with the Operating Environment)• HP Instant Information• HP Technical Documentation web site http://docs.hp.com
Administer an HP-UX system	<p>For Releases Prior to HP-UX 11i Version 3:</p> <ul style="list-style-type: none">• <i>Managing Systems and Workgroups: A Guide for HP-UX System Administrators</i> <p>For Releases beginning with HP-UX 11i Version 3:</p> <ul style="list-style-type: none">• <i>HP-UX System Administrator's Guide</i> (a multivolume set) <p>Other sources of System Administration Information:</p> <ul style="list-style-type: none">• <i>nPartition Administrator's Guide</i>• <i>HP-UX Virtual Partitions Administrator's Guide</i>• <i>Dynamic Root Disk Administrator's Guide</i> <ul style="list-style-type: none">• <i>Planning Superdome Configurations</i> white paper	<ul style="list-style-type: none">• HP Instant Information• HP Technical Documentation web site http://docs.hp.com <ul style="list-style-type: none">• <i>Planning Superdome Configurations</i>, http://docs.hp.com/hpux/onlinedocs/os/11i/superdome.pdf

HP-UX 11i Release Names and Operating System Version Identifiers

With HP-UX 11i, HP delivers a highly available, secure, and manageable operating system that meets the demands of end-to-end Internet-critical computing. HP-UX 11i supports enterprise, mission-critical, and technical computing environments. HP-UX 11i is available on both HP 9000 systems and Integrity systems.

Each HP-UX 11i release has an associated release name and release identifier. The `uname` command with the `-r` option returns the release identifier. Table 2 shows the releases available for HP-UX 11i.

Table 2 HP-UX 11i Releases

OS Version Identifier	Release Name	Supported Processor Architecture
B.11.11	HP-UX 11i Version 1	HP 9000
B.11.23	HP-UX 11i Version 2	Integrity
B.11.23.0409	HP-UX 11i Version 2 September 2004 Update	HP 9000 and Integrity
B.11.31	HP-UX 11i Version 3	HP 9000 and Integrity
B.11.31.0709	HP-UX 11i Version 3, September 2007 Update	HP 9000 and Integrity
B.11.31.0803	HP-UX 11i Version 3, March 2008 Update	HP 9000 and Integrity
B.11.31.0809	HP-UX 11i Version 3, September 2008 Update	HP 9000 and Integrity
B.11.31.0909	HP-UX 11i Version 3, September 2009 Update	HP 9000 and Integrity

For information on supported systems and processor architecture for different versions of HP-UX, refer to the HP-UX system release notes specific to your version of HP-UX (for example, the *HP-UX 11i Version 3 Release Notes*).

Determining Your System Version

The `uname`, `model`, and `swlist` commands can help you determine information about your system, including its hardware type, machine model, operating system version, and operating environment update status. (See `uname(1)`, `model(1)`, and `swlist(1M)`.)

For OS naming conventions, please see “HP-UX 11i Release Names and Operating System Version Identifiers” (page 13).

Table 3 OS Version, System Architecture, and Machine Model

Topic	Command	Sample Output
OS Version	<code>\$uname -r</code>	B.11.31 ¹
Architecture	<code>\$uname -m</code>	ia64 ² 9000/800 ²
Machine Model	<code>\$model</code> ³	ia64 hp server rx5670 9000/800/S16K-A

Table 3 OS Version, System Architecture, and Machine Model (continued)

Topic	Command	Sample Output
Operating Environment	\$ swlist HPUX*OE*	# HPUX11i-OE-MC B.11.31 HP-UX Mission Critical Operating Environment ¹
OS Version.Update	\$ swlist HPUX*OE*	# HPUX11i-TCOE B.11.23.0409 HP-UX Technical Computing OE Component ¹

1 HP-UX 11i OS version identifiers have the form B.11.23 or B.11.23.0409, where B.11.23 is the OS version and 0409 is the year-month of the operating environment (OE) update.

2 ia64 = Integrity. All others = HP 9000.

3 The `getconf MACHINE_MODEL` command gives the same output (see `getconf(1)`).

Typographic Conventions

This document uses the following typographic conventions.

<i>audit</i> (5)	An HP-UX manpage. <i>audit</i> is the name and 5 is the section in the <i>HP-UX Reference</i> . On the web and on the Instant Information DVD, it may be a hot link to the manpage itself. From the HP-UX command line, you can enter <code>man audit</code> or <code>man 5 audit</code> to view the manpage. See <i>man(1)</i> .
<i>Document Title</i>	The title of a document. On the web and on the Instant Information DVD, it may be a hot link to the document itself.
<i>Command</i>	A command name or qualified command phrase.
<i>ComputerOut</i>	Text displayed by the computer.
<i>Emphasis</i>	Text that is emphasized.
Emphasis	Text that is strongly emphasized.
KeyCap	The name of a keyboard key. Note that Return and Enter both refer to the same key.
FirstTerm	The defined use of an important word or phrase.
UserInput	Commands and other text that the user types.
<i>Replaceable</i>	The name of a variable that you may replace in a command or function or information in a display that represents several possible values.
\$	Default user command prompt.
#	Default superuser (<code>root</code>) command prompt.
>	Default continuation command prompt.
\	Line continuation marker.

Examples and Shells

This document describes practices used by the system administrator. Since the `root` user (a superuser) is required to use the POSIX shell /sbin/sh, all command examples use that shell. The POSIX shell is defined in *sh-posix(1)*. For information on other shells, see the *Shells User's Guide* and *sh(1)*.

Command Syntax

Literal	A word or character that you enter literally.
<i>Replaceable</i>	A word or phrase that you replace with an appropriate value.
<code>-chars</code>	One or more grouped command options, such as <code>-ikx</code> . The <i>chars</i> are usually a string of literal characters that each represent a specific option. For example, the entry <code>-ikx</code> is equivalent to the individual options <code>-i</code> , <code>-k</code> , and <code>-x</code> . The plus character (+) is sometimes used as an option prefix.
<code>-word</code>	A single command option, such as <code>-help</code> . The <i>word</i> is a literal keyword. The difference from <code>-chars</code> is usually obvious and is clarified in an Options description. The plus character (+) and the double hyphen (- -) are sometimes used as option prefixes.
<code>[arg]</code>	The bracket metacharacters enclose optional content in formats and command descriptions.
<code>{arg}</code>	The brace metacharacters enclose required content in formats and command descriptions.
<code> </code>	The bar metacharacter separates alternatives in a list of choices, usually in brackets or braces.
<code>arg ...</code> <code>[arg]...</code> <code>{arg}...</code>	The ellipsis metacharacter after a token or a right bracket or a right brace metacharacter indicates that the preceding term and its preceding whitespace, if any, may be repeated an arbitrary number of times.
<code>...</code>	Ellipsis is sometimes used to indicate omitted items in a range.

Function Synopsis and Syntax

HP-UX functions are described in a definition format rather than a usage format. The definition format includes typing information that is omitted when the function call is actually included in a program.

The general definition form is:

```
type func ( type param [, type param]... );
```

For example:

```
int setuname ( const char *name , size_t namelen);
```

The usage form is:

```
func ( param [, param]... );
```

For example:

```
setuname ( name , namelen );
```

The function syntax elements are the same as for commands, except for the options.

Publication History

The document publication date and part number indicate its current edition. The publication date will change when a new edition is released.

To ensure that you receive the new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

First Edition	February 2007 HP Part Number 5991-6480 HP-UX 11i Version 3 Print, Instant Information DVD, and HP Technical Documentation web site http://docs.hp.com
Second Edition	March 2008, HP Part Number 5992-3386, HP-UX 11i Version 3 (B.11.31 March 2008 Update) Print, Instant Information DVD, and HP Technical Documentation web site (http://www.docs.hp.com)
Third Edition	September 2008, 5992-4607, HP-UX 11i version 3 (B.11.31 September 2008 Update) Printed, DVD (Instant Information), and Web (http://www.docs.hp.com)
Fourth Edition	September 2009, 5992-6579, HP-UX 11i version 3 (B.11.31 September 2009 Update) Printed, DVD (Instant Information), and Web (http://www.docs.hp.com)



NOTE: The volumes in the *HP-UX System Administrator's Guide* may be updated independently. Therefore, the latest versions of the volumes in the set may vary with time, with respect to each other. The latest versions of each volume are available at <http://docs.hp.com>.

1 Introduction

The *HP-UX System Administrator's Guide: Configuration Management* continues where the *HP-UX 11i V3 Installation and Update Guide* leaves off. It describes the process of preparing your system for use and connecting it into the local network. Topics include:

- Initial start-up of a preinstalled system (called **instant ignition**). See “Starting A Preloaded System” (page 35).
- Initial configuration of system parameters. See “Setting System and Network Parameters” (page 41).
- Configuring the root user. See “Configuring Root” (page 47).
- Installing and managing user accounts. See “Configuring Users and Groups” (page 49).
- Configuring the system into a network. See “Configuring Networking” (page 75).
- Configuration tools: Systems Insight Manager (HP SIM), System Management Homepage (HP SMH), and HP-UX commands. See “Configuration Tools” (page 19).
- Modifying kernel tunables and modules. See “Configuring the Kernel” (page 153).
- Setting up local and remote printers. See “Configuring Printers” (page 111).
- Organizing e-mail processes. See “Configuring Mail” (page 121)
- Preparing the manpage database. See “Setting Up the Online Manpages” (page 131).

Configuration Tools

There are three ways (at least) to configure your HP-UX system:

- Via the web with the HP Systems Insight Manager (HP SIM).
- Directly on the system or via the web with the text- and web-based HP-UX System Management Homepage (HP SMH).
- The traditional way with HP-UX commands.

HP Systems Insight Manager

The HP Systems Insight Manager (HP SIM) is a major advance in web-based multiple system management. You can use it to configure and manage not only HP-UX systems but also HP Linux systems and HP Windows® systems, as well as non-HP Linux and Windows systems. HP SIM makes use of the HP SMH interface and other tools. It is fully described in the *HP Systems Insight Manager 5.2 Installation and Configuration Guide for HP-UX* on <http://docs.hp.com> (click “Network and Systems Management” and “HP Systems Insight Manager”).

HP-UX System Management Homepage (Web-Based)

The HP-UX System Management Homepage (HP SMH) is a web-based interface that manages one HP-UX system at a time. It completes the transfer of the System Administration Manager (SAM) functions into web-based processes. In HP-UX 11i Version 2, the SAM functional areas for kernel configuration, peripheral devices, and partition management became web-based, using `kcweb`, `pdweb`, and `parmgr`, respectively. In HP-UX 11i Version 3, the full management interface is web-based using HP SMH. For more information, see “[HP-UX System Management Homepage \(Web-Based\)](#)” (page 21). HP SMH is also described in *Simplifying single-system management on HP-UX 11i – HP System Management Homepage (HP SMH)*.

HP-UX System Management Homepage (Text-Based)

The text-based HP-UX System Management Homepage (HP SMH) replaces the System Administration Manager (SAM). It operates as a logged-in management tool. Its principal interface provides simple terminal-based screens. In some cases, the earlier screen-based and X-Window-based displays have been retained in the functional areas. For more information, see “[HP-UX System Management Homepage \(Text-Based\)](#)” (page 25).

HP-UX Commands

HP-UX commands provide you with the finest granularity of access to system configuration, but they require you to be careful to complete all the steps correctly. In some cases, such as root home directory reconfiguration (see “[Create the Root Home Directory](#)” (page 47)), they may be the only way to do it. In other cases, such as tweaking nonpassword data in the password file, it’s just easier to run a command such as `vipw`. Generally, HP recommends that you use HP SMH or HP SIM to manage your system wherever possible and convenient, because they are designed to ensure that all the files and databases are updated and all the commands are issued correctly.



CAUTION: You must use either the command sequence or HP SMH to perform any operation that HP SMH supports. Attempting to start an operation with commands and completing it with HP SMH can result in errors and possibly corrupt data or data structures.

Using HP SIM and HP SMH versus HP-UX Commands

HP SIM and HP SMH hide the complexity of most administrative tasks. They minimize or eliminate the need for detailed knowledge of many administrative commands, thus saving valuable time. Use these utilities whenever possible, especially when first mastering a task. Some tasks described in this document cannot be done by these utilities, in which case you will need to use the HP-UX commands. However, these utilities are the tools of choice for much administrative work.

Many HP SMH procedures provide the option to preview the HP-UX commands that they will use to perform the task at hand.

HP-UX System Management Homepage (Web-Based)

The HP-UX System Management Homepage (HP SMH) helps you with detailed system administration. When you log in via the web (see “[Starting an HP SMH Web Client](#)” (page 24)), you see a **Home** page that gives you access to a wide range of system information and some principal areas of system configuration.

The **Tools** page, selectable from the menu bar, displays the full range of configuration tools available through HP SMH. Depending on which Operating Environment you have installed, your, installation choices, and added optional software, your **Tools** page may be different.

- Accounts for Users and Groups
 - Configure Groups
 - Configure Local Users *or* Configure NIS Users
 - Configure User Templates
- See “Configuring Users and Groups in *HP-UX System Administrator’s Guide: Configuration Management*.
(Also accessible with `/usr/sbin/ugweb`.)
- Audit Configuration
 - Audit Events
 - Audit System Calls
 - Audit Users
- See “Audit Administration” in *HP-UX System Administrator’s Guide: Security Management*.
- Auditing and Security Attributes Configuration(web-based)
 - Audit Events
 - Audit System Calls
 - Local Users
 - System Defaults
- See “Audit Administration” in *HP-UX System Administrator’s Guide: Security Management*.
- Authenticated Commands (PAM)
 - Configure Account Authentication
 - Configure Password Authentication
 - Configure Session Authentication
 - Configure User Authentication

See "Authenticating Users with PAM" in *HP-UX System Administrator's Guide: Security Management*.

- Disks and File Systems

Disks

File Systems

Logical Volumes

Paging Space

Volume Groups

See *HP-UX System Administrator's Guide: Overview*, *HP-UX System Administrator's Guide: Logical Volume Management*, and *VERITAS File System 3.5 (HP OnlineJFS/JFS 3.5) Administrator's Guide* manuals, and *The Next Generation Mass Storage Stack* white paper. The white paper is posted in the Network and Systems Management section of the HP technical documentation web site under Storage Area Management (<http://docs.hp.com/en/netsys.html#Storage%20Area%20Management>).

(Also accessible with /usr/sbin/fsweb)

- Display

Monitor Configuration

X Server Configuration

- Distributed Systems Administration Utilities (DSAU)

Configure Configuration Synchronization

Configure Consolidated Logging

View System Logs

See the *Distributed Systems Administration Utilities User's Guide*.

- Error Management Technology

Query or Customize Error Data

- Event Monitoring Service(web-based)

Requests

- Evweb

Subscription Administration

(Also accessible with /opt/sfm/bin/evweb)

- IPMI Event Viewer

Event Viewer

- Kernel Configuration

Alarms
Modules
Tunables

See “Configuring the Kernel in the *HP-UX System Administrator’s Guide: Configuration Management*.

(Also accessible with `/usr/sbin/kcweb`)

- Network Interfaces Configuration
 - Auto Port Aggregation
 - Network Interface Cards
 - Tunnels
 - Virtual LANs
 - Network Services Configuration
 - Bootable Devices
 - DHCPv6
 - DNS
 - Hosts
 - Network Services
 - NIS
 - Route
 - Networked File Systems
 - System Access
 - Time
 - nPartition Management
 - View and Manage Complex
 - View and Manage Remote Complex
 - Peripheral Devices
 - Manage Peripheral Devices
 - OLRAD Cards
 - I/O Tree
- See “Configuring Peripherals” in the *HP-UX System Administrator’s Guide: Configuration Management*.
- (Also accessible with `/usr/sbin/pdweb`)
- Printer Management
 - Configure Printers or Plotters
 - Manage Print Requests
 - Save or Restore Print Spooler Configuration
- See “Configuring Printers” in the *HP-UX System Administrator’s Guide: Configuration Management*.

- Printer Management(web-based)
 - Configure Printers or Plotters
 - Manage Print Requests
 - Save or Restore Print Spooler Configuration

See “Configuring Printers” in the *HP-UX System Administrator’s Guide: Configuration Management*.
- Resource Management
 - Event Monitoring Service
- Serviceguard
 - Serviceguard Manager

See the *Managing Serviceguard* manual.
- Software Management
 - Copy Depot Software
 - Install Software
 - List Depot Software
 - List Installed Software
 - Remove Depot Software
 - Remove Installed Software

Over time with new OEs the tools available in HP SMH may increase allowing you to perform more system administration tasks via HP SMH

See *HP-UX System Administrator’s Guide: Routine Management Tasks and Software Distributor Administration Guide*.

Starting an HP SMH Web Client

A majority of the applications in web-based HP SMH support multi-byte locales.

Supported browsers are specified in the *HP System Management Homepage Release Notes*, available on <http://docs.hp.com>.

You should have an X Window System client running on the local system. See “Using the X Window System” (page 32) for details.

In a browser window, enter one of the following URLs, where *targetsystem* is the host domain name or IP address of the system being managed:

<http://targetsystem:2301/> You will see a web page with the message:

Please Wait... You will be
redirected to the System Management
Homepage

If the hpsmh daemon has not been started on the target system, it is automatically started.

After a brief wait, you will see the login web page for HP SMH. Enter a valid HP SMH login name (for example, root) and the password.

<https://targetsystem:2381/>

- If the hpsmh daemon is running on the target system, you will see the login web page for HP SMH. Enter a valid HP SMH login name (for example, root) and the password.
- If the hpsmh daemon has not been started on the target system, you will see a browser-dependent message, such as:
 - The page cannot be displayed (Internet Explorer)
 - Unable to connect (Firefox)
 - The document contains no data (Mozilla)Try using the URL for port 2301 (above).

Starting the HP SMH Daemon on the Target System

To start the HP SMH daemon on the target system, enter the following command on the target system:

```
# /opt/hpsmh/bin/hpsmh autostart
```

The HP SMH daemon is also started if a browser attempts to access port 2301 on the target system, as with the URL: <http://targetsystem:2301/>

HP-UX System Management Homepage (Text-Based)

The HP-UX System Management Homepage (HP SMH) helps you with the administration of your system.

HP SMH administrative areas:

- a - Auditing and Security
 - Audited Events
 - Audited System Calls
 - Audited Users
 - System Security Policies
 - Authenticated Commands

See the *HP-UX System Administrator's Guide: Security Management*.

- c - Security Attributes Configuration
Topics include password aging and default umask.
 - s - System Defaults
Configure system-wide values of security attributes.
See *security(4)*.
 - l - Local Users
Configure per-user values of security attributes of local users (if NIS is *not* configured).
See *userdb(4)*.
- See "Configuring System and User Security" in the *HP-UX System Administrator's Guide: Configuration Management*.
- (Also accessible with `/usr/sbin/secweb -t`)
- d - Peripheral Devices
 - o - OLRAD Cards
View all OLRAD-capable slots and cards on the system and perform OL* operations
See *Interface Card OL* Support Guide*.
 - i - I/O Tree
View all devices on the system
- (Also accessible with `/usr/sbin/pdweb -t`)
- e - Resource Management
 - Event Monitoring Service
See *Using the Event Monitoring Service*.
(Also accessible with `/opt/sfm/bin/evweb`)
 - f - Disks and File Systems
 - f - File Systems
View or Manage File Systems
 - l - Logical Volumes
View or Manage Logical Volumes

- v - Volume Groups
View or Manage Volume Groups
- d - Disks
View or Manage Disk Devices

See *HP-UX System Administrator's Guide: Overview* and *HP-UX System Administrator's Guide: Logical Volume Management*.

(Also accessible with /usr/sbin/fsweb)

- g - Display
 - Monitor Configuration
 - Xserver Configuration
- k - Kernel Configuration
 - t - Tunables

View or modify kernel tunables

See *kctune(1M)*
 - m - Modules

View or modify kernel modules and drivers

See *kcmodule(1M)*
 - a - Alarms

View or modify alarms for kernel tunables

See *kcalarm(1M)*
 - l - Log Viewer

View the changes made to kernel tunables or modules

See *kclog(1M)*
 - u - Usage

View usage of kernel tunables

See *kcusage(1M)*
- c - Manage Configuration

View the options available to manage configurations

See *kconfig(1M)*
- b - Restore Previous Boot Values

Restores Previous Boot Values for Tunables and Modules

See "Configuring the Kernel" in the *HP-UX System Administrator's Guide: Configuration Management*.

(Also accessible with /usr/sbin/kcweb -t)

- l - Printers and Plotters(new)
 - p - printers
 - r - print requests
 - s - spooler configuration
 - a - add printer
 - c - show common problems
- m - Event Monitoring Service
 - q - requests
 - r - resources
- n - Networking and Communications
 - s - Network Services Configuration
 - b - Bootable Devices
 - b - DHCP Device Groups Booting from This Server
 - r - Devices for which Boot Requests are Relayed to Remote Servers
 - f - Fixed-Address Devices Booting from This Server
 - v - DHCPv6
 - c - Configuring DHCPv6 Server
 - s - Configuring Default DHCPv6 Client Settings
 - h - Configuring a Host to Act as a DHCPv6 Relay Agent
 - r - Configuring DHCPv6 Relay Interface Mappings
 - p - Configuring DHCPv6 Address Pools
 - d - Configuring DHCPv6 Client Duid Groups
 - g - Configuring DHCPv6 Device Groups
 - d - DNS (BIND)
 - 1 - DNS Local Name Server
 - r - DNS Resolver
 - h - Hosts
 - h - Local Hosts File
 - n - NIS
 - s - Name Service Switch
 - k - Network Services
 - f - Networked File Systems
 - s - Share/Unshare File System
 - a - Automounted Remote File Systems

- n - Netgroups
 - Local Netgroups
 - o r - Routes
 - o c - System Access
 - i - Internet Services
 - r - Remote Logins
 - o t - Time
 - s - System Clock
 - b - NTP Broadcasting
 - n - NTP Network Time Sources

See “Configuring Networking” in the *HP-UX System Administrator’s Guide: Configuration Management*.

See also *NFS Services Administrator’s Guide*.

(Also accessible with `/usr/sbin/ncweb -t`)

- i - Network Interfaces Configuration
 - o a - Auto Port Aggregation
 - o n - Network Interface Cards

See *HP-UX LAN Administrator’s Guide*.

 - o v - Virtual LANs

See *HP-UX VLAN Administrator’s Guide*.

 - o t - Tunnels

(Also accessible with `/usr/sbin/ncweb -t`)

- p - Printers and Plotters
 - Print Requests
 - Printers and Plotters
 - Save/Restore Spooler Configuration

See “Configuring Printers” in the *HP-UX System Administrator’s Guide: Configuration Management*.

- s - Software Management
 - i - Install Software
 - r - Remove Installed Software
 - l - Interactive List, Installed Software
 - s - Quick List, Installed Software
 - p - Quick List, Installed Patches
 - c - Copy Depot Software

- m - Remove Depot Software
 - d - Interactive List, Depot Software
 - u - Update HP-UX Operating Environment
 - u - Accounts for Users and Groups
 - l - Local Users

View or Configure Local Users (if NIS is *not* configured).
 - g - Groups

View or Configure Groups.
 - t - Templates

View or Configure User Templates.
- See “Configuring Users and Groups” in the *HP-UX System Administrator’s Guide: Configuration Management*.
- (Also accessible with `/usr/sbin/ugweb -t`)

Starting Text-Based HP SMH

Text-based HP SMH only supports the C (English) locale. HP recommends that you set your locale variables, such as `LANG` and `LC_ALL`, to C.

To run HP SMH, you must be superuser or have been granted access (see “Giving Users Limited Access to Text-Based HP SMH” (page 31)).

Portions of HP SMH can use the X Window System to display enhanced screens. You can choose to have those screens displayed as text graphics instead. The X screens allow you to use the mouse pointer to navigate the screens. The text screens and the menu displays use keyboard controls, notably **Tab**, the arrow keys, **Enter**, and certain letter keys, as indicated on the screen. **Esc** usually ends the current operation, going to the previous screen. On a menu display, **X** terminates the program.

- To start text-based HP SMH modules with the X Window interface,
 1. Enable the X Window system as described in “Using the X Window System” (page 32).
 2. Run the module with the `-t` option:

```
# /usr/sbin/fsweb -t      # Disks and File Systems
# /usr/sbin/kcweb -t      # Kernel Configuration
# /usr/sbin/ncweb -t      # Networking and Communications
# /usr/sbin/pdweb -t      # Peripheral Devices
# /usr/sbin/secweb -t     # Security Attributes Config
# /usr/sbin/ugweb -t      # Accounts for Users and Groups
```



NOTE: The -t option is not available for /usr/sbin/smh.

If the X Window interface is not available, the modules use the alternate text graphic display.

- To start text-based HP SMH without the X Window interface, you must unset the DISPLAY environment variable. For example, you can enter the following:

```
# ( unset DISPLAY ; /usr/sbin/smh )
```

This unsets the DISPLAY variable while HP SMH executes. When HP SMH ends, the value of DISPLAY is restored. Notice the enclosing parentheses and the semicolon between the commands.

Similarly, you can start the modules without the X Window interface. For example,

```
# ( unset DISPLAY ; /usr/sbin/ugweb )
```

Giving Users Limited Access to Text-Based HP SMH

As system administrator, you can give limited text-based HP SMH access to nonsuperusers individually by user name and collectively by primary group name.



NOTE: The privileges assigned to users and groups by the text-based restricted HP SMH do not apply to the web-based HP SMH.

1. Activate Restricted HP SMH.

```
# /usr/sbin/smh -r
```

2. You can assign text-based HP SMH privileges by user and by group. You can toggle between the lists of defined users and groups with the **u** and **g** keys, respectively.

3. To select a user or group, move the highlight to that entry and press **Enter**. The list of text-based HP SMH areas is displayed.

```
Resource Manager
Disks and File Systems
Display
Kernel Configuration
Printers and Plotters
Networking and Communications
Peripheral Devices
Security Attributes Configuration
Software Management
Auditing and Security
Accounts for Users and Groups
```

4. Choose one of the following:

- To assign an area, highlight it and press **e**.
- To assign all areas, press **E**.

- To disable an area, highlight it and press **d**.
- To disable all areas, press **D**.

You can repeat these operations in any combination. The changes are displayed each time you press a key.

5. Press **s** to save the changes.
6. Press **Esc** to return to the previous screen.
7. Press **x** to exit from the program.

User and group privileges are managed separately. Group privileges apply to all users for which it is their primary group, as shown in `/etc/passwd`. Users can acquire a privilege individually, through their groups, or both.

When privileged users run `/usr/sbin/smh`, they run text-based HP SMH. They have superuser status in the defined areas and will only see those HP SMH areas in the menu. All other areas of HP SMH are hidden.

When users without special access to HP SMH try to run text-based `smh`, they receive a message like the following:

```
Neither the user 'allanp' nor his primary group 'users' has
Restricted sam privileges. Exiting!
```

Using the X Window System

Web-based HP SMH has a few commands and displays that require that your local system (the client) and the target system be running an X Window System™ client. On an HP-UX system (client or target), this is provided by the `X11-RUN` fileset, which is normally installed. On a non-HP-UX client system, such as one running Microsoft® Windows®, you may need to install a third-party client program.

Text-based HP SMH has text alternatives to the X Window displays if `DISPLAY` is not set or there is no active X Window client.

Setup for an HP-UX Local System

If you are logged in to the target system from an HP-UX client system, you need to do the following:

- On your HP-UX client system, enable it to receive X Window screens from the target system:


```
# /usr/bin/X11/xhost + targetsystem
```

The `xhost` command is described in the `xhost(1)` manpage.
- On the target system, set your `DISPLAY` environment variable to the domain name or IP address of your client system:


```
# export DISPLAY=clientsystem:0.0
```

Setup for a Non-HP-UX Client System

If you are logged in to the target system from a non-HP-UX client system, you need to do the following:

- On your non-HP-UX client system, start your X Window server and determine your IP address. This is often available in the **Help About** menu item of the program screen. If the client is connected to an HP-UX system, the command echo \$DISPLAY will show the IP address.
- On the target system, set your DISPLAY environment variable to the IP address of your client system:

```
# export DISPLAY=clientsystemIP:0.0
```

2 Configuring System Parameters

This chapter describes the process of starting a preloaded system and setting general system parameters. The topics include:

- “Starting A Preloaded System” (page 35)
- “Preventing Users from Logging In” (page 36)
- “Changing Login Special Characters” (page 37)
- “Controlling Usage and Processes with Run-Levels” (page 38)
- “Setting the System Clock” (page 40)
- “Setting System and Network Parameters” (page 41)
- “Customizing System-Wide and User Login Environments” (page 42)
- “Setting Shadow Password Mode” (page 43)
- “Setting Long User, Group, Host, and File Names” (page 43)
- “Configuring /etc/hosts” (page 44)
- “Configuring New HP-UX Systems into Workgroups” (page 45)
- “Configuring a New System into a Network” (page 45)

Starting A Preloaded System

System administrators can use these directions as a quick reference or print them out for users about to start up their own systems.



IMPORTANT: System security is an important part of system configuration. HP-UX provides a wide variety of security features, including basic file and access control, Trusted System configuration, intrusion detection with HP-UX HIDS, and system “lockdown” with HP-UX Bastille. Use the *HP-UX System Administrator’s Guide: Security Management* to develop a security plan that meets your needs. You can install and configure that plan as part of the following steps.

1. Turn on the monitor and computer system.

The system will run a series of self-tests. For information about these self-tests, see your *Owner’s Guide*.

After a short time, a series of messages is displayed as various hardware and software subsystems are activated. Unless something is wrong, you are not asked to respond to these messages.

2. Enter information as it is requested.

You will need to know your host name and IP address. Your network administrator can provide you with the host name and IP address.

Press **Return** to use the default values. To provide missing information later, log in to a terminal as superuser and execute the command `/sbin/set_parms`. See “Setting System and Network Parameters” (page 41).

3. Specify a root password.

The user name for the superuser is `root`.

The system completes its start-up sequence and displays the desktop login screen.

4. Log in as `root` for your first session.
5. Establish the environment for the `root` user. See “Configuring Root” (page 47).
6. Set up and configure additional security, as suggested in the **Important** note above. See the *HP-UX System Administrator’s Guide: Security Management*. Some security measures might have been set up during the installation process. See the *HP-UX 11i v3 Installation and Update Guide*.
7. Add users as needed. See “Configuring Users and Groups” (page 49).
8. Set up NFS if desired. See “Configuring Networking” (page 75).

Preventing Users from Logging In

One of the issues for system administrators is how to keep nonsuperusers from logging in to a system while system configuration or system maintenance is underway. This can be accomplished with the combination of the `/etc/nologin` file and the `NOLOGIN` variable in the `/etc/default/security` file.

If `/etc/nologin` exists and `NOLOGIN=1`, nonsuperusers are barred from logging in. What they will see is a system-generated message (`Only...security`) followed by the contents of `/etc/nologin`. For example:

```
Only superusers are allowed to login at this time due to the presence  
of the file /etc/nologin and NOLOGIN option set in /etc/default/security
```

```
--+ SYSTEM MAINTENANCE +-  
This system (hprdc185) is undergoing system maintenance. We expect to  
return to full service on Thursday morning at 9 a.m. Pacific Time.  
Thank you.
```

Procedure 2-2 To set up the login barricade

1. Create or edit the file `/etc/nologin`. It can be empty, but users might appreciate some information on why the system is unavailable.
2. Edit the file `/etc/default/security` and set the `NOLOGIN` line to `NOLOGIN=1`.

Procedure 2-3 To restore normal system access

- Do at least one of the following:
 - a. Edit the file `/etc/default/security` and set the `NOLOGIN` line to `NOLOGIN=0`.
 - b. Delete the file `/etc/nologin`.

You can also use the `ch_rc` command to modify the `NOLOGIN` variable (see *ch_rc(1M)*).

```
# /usr/sbin/ch_rc -a -p NOLOGIN=1 /etc/default/security
# /usr/sbin/ch_rc -a -p NOLOGIN=0 /etc/default/security
```

Changing Login Special Characters

As many users discover to their frustration, UNIX was originally developed on TeleType machines that had no backspace operation. Consequently, the process of editing input data was performed by printing characters: #, delete one character and @, delete the line. Interrupt was defined as Rubout (ASCII DEL). This situation has carried over to the process of logging in to a system. #, @, and DEL are still the HP-UX default login control characters today.

Once they have logged in, users get around this inconvenience by including the `stty` command in their `.profile` and `.login` scripts to change DEL, #, and @ to commonly used control characters, such as ETX (**Ctrl-C**), backspace (**Ctrl-H**), and NAK (**Ctrl-U**), respectively.

You can change the login control characters used by such programs as `getty`, `rlogin`, `ssh`, and `telnet`. To do this, you use `stty` to set the desired parameters on a special device file named `/dev/ttyconf`. See `stty(1)` for details.

The `/dev/ttyconf` file provides a way to change the default behavior for all logins by all users. `/dev/ttyconf` represents a terminal device and maintains a set of all the terminal control characters that can be displayed by `stty`. The parameters of any terminal device file can be viewed and modified by passing the file as input to the `stty` command.

When a system is rebooted, the contents of `/dev/ttyconf` are reset to the default values, as displayed here by `stty` with the `-a` option:

```
# stty -a < /dev/ttyconf
min = 4; time = 0;
intr = DEL; quit = ^\; erase = #; kill = @@
eof = ^D; eol = ^@; eol2 <undef>; swtch <undef>
stop = ^S; start = ^Q; susp <undef>; dsusp <undef>
werase <undef>; lnext <undef>
```

To change any of the displayed parameters, execute `stty` with the appropriate options. For example, to change the system defaults for `intr`, `erase`, and `kill` to **Ctrl-C**, **Ctrl-H**, and **Ctrl-U**, respectively, enter the command:

```
# stty intr ^C erase ^H kill ^U < /dev/ttyconf
```

To see just the variations from the defaults, enter `stty` without options:

```
# stty < /dev/ttyconf
intr = ^C; erase = ^H; kill = ^U;
swtch <undef>;
```

Unlike the other login commands, the `getty` command does not automatically use `/dev/ttyconf`; it requires the `-f` option. In the `/etc/inittab` file, add the `-f` option to each `getty` command. The next time the terminal device is reopened, `getty` will

use the new settings. In order to reset the system console, you must reboot the system. See *getty(1M)* for details.

To set */dev/ttyconf* every time the system boots, add the *stty* command to */etc/inittab*. If you place it before the *getty* command for the console, the console will also use the revised control characters:

For example, change:

```
cons:123456:respawn:/usr/sbin/getty console console      # system console
ttpl:234:respawn:/usr/sbin/getty -h tty0p1 9600
```

To:

```
ttco::bootwait:/sbin/stty intr ^C erase ^H kill ^U < /dev/ttyconf
cons:123456:respawn:/usr/sbin/getty -f console console  # system console
ttpl:234:respawn:/usr/sbin/getty -f -h tty0p1 9600
```

Controlling Usage and Processes with Run-Levels

A **run-level** is an HP-UX state of operation in which a specific set of processes is permitted to run. These processes and default run-levels are defined in the file */etc/inittab*.

The run-levels are:

- Run-level s The operating mode system administrators use (often called **single-user mode**). This mode ensures that no one else is on the system while you are performing system maintenance tasks. In this run-level, the only access to the system is through the system console by the user root. The only processes running on the system can be the shell on the system console, background daemon processes started by */sbin/rc*, and processes that you invoke. Commands requiring an inactive system (such as */sbin/fsck*) should be run in run-level s.
- Run-level 1 Starts a subset of essential system processes; can also be used to perform system administrative tasks.
- Run-level 2 The operating mode typically called **multiuser mode**. This mode allows all users to access the system.
- Run-level 3 For NFS servers. In this mode, NFS file systems can be shared, as required for NFS servers.
For CDE users. In this mode, CDE is active. CDE is the default desktop on HP-UX 10.30 and later.
- Run-level 4 Sometimes used by optional software.
Depending on the software installed on your system, the default run-level is usually run-level 3 or 4. The default run-level for CDE is 3.

To determine the current run-level of the *init* process, type:

```
# who -r
. . . run-level 3 Mar 5 12:01 3 0 S
```

You can add to and change the sequence of processes that HP-UX starts at each run-level. For more information, see the *HP-UX System Administrator's Guide: Overview*. Also see the *inittab(4)* manpage.

A superuser logged in at the system console can also change the current run-level with the /sbin/init and /sbin/shutdown commands, as follows:

1. Warn all users who are currently logged in. Whenever the run-level of the system is changed, any process that does not have a run-level entry matching the new run-level will be killed. There is a grace period of 20 seconds after an automatic warning signal is sent.
2. To change to run-level s, use the shutdown command.
To change to a run-level other than run-level s, use the init command.
See *shutdown(1M)* and *init(1M)*.



CAUTION: Only use the shutdown command to change to run-level s (that is, do not use /sbin/init s).

The shutdown command *safely* brings your system to run-level s without leaving system resources in an unusable state. The shutdown command also allows you to specify a grace period to allow users to terminate their work before the system goes down. For example, to enter run-level s after allowing 30 seconds, enter:

```
# shutdown 30
```

To shut down immediately, enter one of the following:

```
# shutdown now
# shutdown 0
```

To achieve a true single-user mode with a quiet system, the best tactic is to reboot the system with an interrupted boot. See *hpxx(1M)* and *hpxx.efi(1M)* and the *HP-UX System Administrator's Guide: Routine Tasks*.

Do not use run-level 0; this is a special run-level reserved for system installation.

For increased security, ensure that the permissions (and ownership) of the files /sbin/init and /etc/inittab are as follows:

```
# ll /sbin/init /etc/inittab
-r--r--r-- 1 root      sys          2152 Oct 17 01:25 /etc/inittab
-r-xr-xr-x 1 bin      bin         1968452 Oct 10 21:31 /sbin/init
```

Also, the size of inittab should be just a few KB while the size of init should be one to two MB, as shown.

Setting the System Clock

Only a superuser (`root`) can change the system clock. The system clock budgets process time and tracks file access.

Potential Problems When Changing the System Clock

The following are potential problems you can cause by changing the system clock:

- The `make` program is sensitive to a file's time and date information and to the current value of the system clock. Setting the clock forward will have no effect, but setting the clock backward by even a small amount may cause `make` to behave unpredictably.
- Incremental backups depend heavily on a correct date because the backups rely on a dated file. If the date is not correct, an incorrect version of a file can be backed up.
- Altering the system clock can cause unexpected results for jobs scheduled by `/usr/sbin/cron`:
 - If you set the clock ahead, `cron` attempts to catch up by immediately starting all jobs scheduled to run between the old time and the new. For example, if you set the clock ahead from 9:00 to 10:00, `cron` immediately starts all jobs scheduled to run between 9:00 and 10:00.
 - If you set the time back, `cron` does not run any jobs until the clock catches up to the point from which it was set back. For example, if you set the clock back from 8:00 to 7:30, `cron` will not run any jobs until the clock again reaches 8:00.
 - If you set the time back just after `cron` starts a job but before the job is recorded, the job will be run twice. For example, if a job scheduled for 8:00 is started and the time is set back to 7:30 before the job is recorded, the job will be recorded as starting at about 7:30. When the clock again reaches 8:00, `cron` will start the job a second time.

Setting the Time Zone (TZ)

To change the local time zone, you can use the `/sbin/set_parms timezone` command. See “[Setting System and Network Parameters](#)” (page 41). This change requires a system reboot.

Setting the time zone only affects how time is converted to local time for display. Internally, the system records time in Universal Time (UTC).

Setting the Time and Date

If you have to reset the time or date, you can use the `set_parms date_time` command or the `date` command. See “[Setting System and Network Parameters](#)” (page 41), `set_parms(1M)`, and `date(1)`.



NOTE: HP strongly recommends that you use single-user mode when changing the system clock. Therefore, warn users of a planned system shutdown. See the *HP-UX System Administrator's Guide: Routine Tasks* for details on system shutdown.



CAUTION: Changing the date while the system is running in multiuser mode may disrupt user-scheduled and time-sensitive programs and processes. Changing the date may cause *make(1)*, *cron(1M)*, and the Source Control subsystems SCCS (*sccs(1)*) and RCS (*rcs(1)*) to behave in unexpected ways. Additionally, any HP or third-party supplied programs that access the system time or the file time stamps stored in the file system, may behave in unexpected ways after the date is changed. *Setting the date backward is not recommended*. If changes were made to files in SCCS file format while the clock was not set correctly, check the modified files with the *val* command. See *val(1)* for details. See "Potential Problems When Changing the System Clock" (page 40) for more information.

Setting System and Network Parameters

The */sbin/set_parms* program is run automatically when you first boot the system after installation. If you do not install HP-UX onto the system yourself, or you do not provide networking information during the installation, you can add this information later by running */sbin/set_parms initial*. See *set_parms(1M)* for details.

You can reset networking parameters at any time by running */sbin/set_parms* again and rebooting the system. Any modifications should be made as soon as possible after the initial installation.

To enter the appropriate *set_parms* dialog screen to manually add or modify information after booting, log in as superuser and specify

```
# set_parms keyword
```

where *keyword* is one of the keywords in Table 2-1. You will be prompted for the appropriate data. The list of keyword choices is displayed when you enter *set_parms* without a keyword:

```
# set_parms
Usage: set_parms <argument>
      Where <argument> can be:
hostname
timezone
date_time
root_passwd
ip_address
addl_netwrk
      or initial (for entire initial boot-time dialog sequence)
```

Changes you make using *set_parms* will take effect after you reboot the system. See the *HP-UX System Administrator's Guide: Routine Tasks*.



NOTE: If a system is having trouble communicating with other systems, check that /etc/rc.config.d/netconf, /var/adm/inetd.sec, and /etc/hosts files all contain the correct official host name.

Table 2-1 The set_parms Keywords

Keyword	Description
initial	Run the entire initial boot-time dialog sequence, in the order hostname, timezone, date_time, root_passwd, ip_address, addl_network.
hostname	Set your unique system or “node” name. This name must contain only alphabetic characters, numbers, underscores, or dashes, and must start with an alphabetic character. The maximum name length is eight characters unless long host names is set, when the maximum is 255 characters. See “Setting Long Host Names” (page 44) for details.
timezone	Set the time zone where your system is located. Changing the time zone does not affect the system clock or file dates, which are always maintained in Universal Time (UT).
date_time	Set the current date and time for the time zone. See “Setting the System Clock” (page 40) for important details.
root_passwd	Set the root password if the current password is null, usually when the system is first initialized. Otherwise, it does nothing.
ip_address	Set the internet protocol (IP) address. If networking is installed, this is an address with four numeric components, each of which is separated by a period with each number between 0 and 255. For example the IP address of example.com is: 192.0.34.166. If you do not have networking installed, you will not be prompted for the IP address.
addl_netwrk	Set additional network parameters. These allow you to configure additional network parameters, such as the subnetwork mask, network gateway, network gateway IP address, local domain name, Domain Name System (DNS) server host name, DNS server IP address and Network Information Service domain name.
locale	Configure local language settings. You can select the language from a menu provided by set_parms or enter a new language (not shown in the menu) by selecting Others. set_parms will verify that a user-specified language is installed. If not, you must install the language before you can use it with set_parms.

Customizing System-Wide and User Login Environments

Defaults for system-wide variables, such as time-zone setting, terminal type, search path, and mail and news notification, can be set in /etc/profile for POSIX and Korn shell users and in /etc/csh.login for C shell users.

User login scripts can be used to override the system defaults. When HP SMH or useradd adds a user, default user login scripts are copied to the user's home directory from the skeleton directory, which defaults to /etc/skel. See "Changing the Skeleton Directory" (page 54). The POSIX and Korn shells use .profile. The C shell uses .login and .cshrc. See the *Shells: User's Guide* and the *Technical Addendum to the Shells: User's Guide* for information on customizing user login scripts.



NOTE: Do a full backup once you have initially set up and customized your system. This allows you to reconstruct your system — kernel, system files, file system structure, user structures, and your customized files — if you need to. Use HP-UX commands to perform the backup, as described in the *HP-UX System Administrator's Guide: Routine Tasks*.

Setting Shadow Password Mode

Shadow password mode is a state in which account and password security information and passwords are stored in a file, /etc/shadow, that can only be accessed by a superuser. The standard password file, /etc/passwd, retains all the other customary information, except that the password fields for all users is changed to an x.

See *pwconv(1M)* and *shadow(4)* for details.

To Switch to Shadow Password Mode

Execute the following command:

```
# pwconv
```

For each entry in /etc/passwd, the password is changed to x, and the password and password aging information are transferred to /etc/shadow.

To Switch to Standard Password Mode

Execute the following command:

```
# pwunconv
```

For each entry in /etc/shadow, the password and the password time limit information are stored in the password field of the user entry in /etc/passwd. Account aging and the password warn limit are discarded. The /etc/shadow file is deleted.

Setting Long User, Group, Host, and File Names

Setting Long User and Group Names

By default, user names are restricted to eight characters and group names to 16 characters. To set long (up to 254-character) user and group names, see *lugadmin(1M)*.

Once long user and group names have been set and used, you should not attempt to revert to short names.

Setting Long Host Names

If the kernel tunable `expanded_node_host_names` is off (0), the maximum host node name length is eight characters (for example, `hprdc185`) and the maximum full host name length is 64 characters (for example, `hprdc185.example.com`). If it is on (1), the maximum for both is 255 characters. By default, this tunable is off. See *expanded_node_host_names(5)* for details.

Setting Long File Names

The `convertfs` command changes an existing file system to long file names. The `newfs` command creates a new file system with short (-S) or long (-L) file names. See the *convertfs(1M)* and *newfs(1M)* manpages and the *HP-UX System Administrator's Guide: Routine Tasks*.

A short name is up to 14 characters; a long name is up to 255 characters. A long file name system cannot be converted back to short file names.



NOTE: The `/usr` directory should be in a long file name system, since many manpage names exceed 14 characters.

Configuring `/etc/hosts`

You can use any text editor to edit the `/etc/hosts` file. If you are not running BIND, you can use HP SMH.

1. If no `/etc/hosts` file exists on your system, copy `/usr/newconfig/etc/hosts` to `/etc/hosts`, or use FTP to copy another system's `/etc/hosts` file to your system. See the *ftp(1)* manpage for more information.

2. Make sure the `/etc/hosts` file contains the following line:

```
127.0.0.1      localhost loopback
```

3. Add your own host's IP address, name, and aliases to the `/etc/hosts` file, as in the following example:

```
15.nn.xx.103      wszx6 patrick
```

The first field is the IP address, the second is the official host name (as returned by the `hostname` command), and any remaining fields are aliases. See the *hosts(4)* manpage.

4. If the system has more than one network card, add a line to `/etc/hosts` for each IP address. The entries for the additional cards should have the same official host name but different aliases and different IP addresses.

5. Add the names of any other hosts that you need to reach. If you will be using a BIND or NIS server on a different host, add the name of that host.

If your site uses DNS (Domain Name Service) or NIS (Network Information Service), `/etc/hosts` acts as a backup resource in case the name server goes down; so it is a good idea to add the names of systems that the local system frequently needs to reach.

Configuring New HP-UX Systems into Workgroups

If you have a group of closely related HP-UX systems, configure the new system into the group by doing the following tasks:

- Set up NFS mounts to allow the system's users to share working directories. See “[Adding a User to Several Systems: A Case Study](#)” (page 70) or “[Sharing Remote Work Directories](#)” (page 69).

If you are using NIS, you can use the `/etc/netgroup` file to define network-wide groups used for permission checking when doing remote mounts, remote logins, and remote shells. See the manpage *netgroup(4)*.

- Configuring NFS mounts. See “[Mounting a Shared File System \(HP-UX to HP-UX\)](#)” (page 78)
- Add local users and groups. See “[Configuring Users and Groups](#)” (page 49).
- Add remote printers. See “[Adding a Remote Printer to the LP Spooler](#)” (page 113).

Configuring a New System into a Network

To configure a new system into the network:

- Set the network information. See “[Setting System and Network Parameters](#)” (page 41).
- Enable network services. See “[Allowing Access from Remote Systems](#)” (page 72).
- Enable X server access. See “[Enabling X Window Server Access](#)” (page 73)
- Set up printers. See “[Configuring Printers](#)” (page 111).
- Add software as needed. See the *HP-UX System Administrator's Guide: Routine Tasks*.

3 Configuring Root

Traditionally, the HP-UX root user, or superuser, has used the system's root directory, /, as its home directory. This means that root's supporting files, such as .profile, .kshrc, and .sh_history, are mixed in with file system mounts and other public data.

HP recommends that you create a separate home directory for root. Since it must be on the root volume, rather than in the /home directory, which is often on a separate file system, we suggest you use a directory name like /homeroott, which would reside on the root volume. By having a private home directory, you also provide a secure location for the root user's private files.

This chapter describes how to configure the home directory and the environment for the root user.

Create the Root Home Directory

You need to use HP-UX commands, rather than HP SMH.

1. Create a home directory for the root user. This directory must reside in the system root directory (/), since it must be available whenever the system is running. A useful directory name is /homeroott. Create it and make it accessible only by root.

```
# mkdir /homeroott  
# chmod 700 /homeroott
```

2. Modify /etc/passwd to make this new directory the home directory for the root user, as follows. The steps are shown in the example below.
 - a. Use the vipw command to load /etc/passwd into vi.
 - b. Insert homeroott after the / in field 6 of line 1.
 - c. Identify the system in comments field 5.
 - d. Save the file and exit.

```
# vipw  
root:3ngTYOinJA.Mc,/0WR:0:3:::/sbin/sh  
    ihomeroottEsc  
root:3ngTYOinJA.Mc,/0WR:0:3:::/homeroott:/sbin/sh  
    iMySystem Root UserEsc  
root:3ngTYOinJA.Mc,/0WR:0:3:MySystem Root User:/homeroott:/sbin/sh  
:wq  
#
```



IMPORTANT: Points of interest:

- The root entry must be the first line of the /etc/passwd file.
 - The user ID in field 3 is 0.
 - Conventionally, the group ID in field 4 is 3 for group sys.
 - The system entry in field 5 can identify which root user on printouts to networked printers.
 - The absolute home directory is in field 6.
 - The login shell in field 7 *must* be /sbin/sh.
-

3. When NIS is configured, run /var/yp/ypmake to build or synchronize the passwd maps.
4. Move any private files into root's home directory. For example, dot files (/ . [A-zA-Z] *).
`# mv /. [A-zA-Z]* /homeroot`
5. Verify that you can log in as root or su - root from a different session. If you can't, you still have this session to make corrections.

4 Configuring Users and Groups

You can control who has access to your system, its files, and its processes.

Authorized users gain access to the system by supplying a valid user name (login name) and password. Each user is defined by an entry in the file `/etc/passwd`. You can use HP SMH to add, remove, deactivate, reactivate, or modify a user account.

For additional information about passwords, refer to *passwd(4)* and *passwd(1)*. To manually change user account entries, use the `/usr/sbin/vipw` command to edit `/etc/passwd`; see *vipw(1M)* for details.

For security information, see the *HP-UX System Administrator's Guide: Security Management*.

You can add a user in several ways:

- “Adding a User with Text-Based HP SMH” (page 50).
- “Manually Adding a User” (page 61).
- “Automating the Process of Adding a User” (page 60).

Consider performing the following tasks for your new user:

- Add a user to a group. See “Defining Group Membership” (page 67).
- Add a user to mail distribution lists.
- Add a user to disk quota systems.
- Allow user to log in from other systems without a password. See “\$HOME/.rhosts File” (page 72).
- Mount remote directories using NFS. See “Configuring the Network File System (NFS)” (page 75).
- Give remote access to a user. See “Allowing Access from Remote Systems” (page 72).
- Set up the user’s login environment. See “Customizing System-Wide and User Login Environments” (page 42).
- Test the new account.

Configuring Users and Groups with HP SMH

To add a user, perform the following tasks:

- Ensure that the user has a unique UID.
- Insert a line for the user in the `/etc/passwd` file.
- Make a home directory for the user.
- Create an environment for the user.

Adding a User with Text-Based HP SMH

Use this procedure to add a new user.

If you are adding a number of users with the same basic characteristics, consider using a template. See “[Making User Templates with Text-Based HP SMH](#)” (page 55) and “[Using a Template to Add a User with Text-Based HP SMH](#)” (page 55).

1. Start HP SMH, as described in “[Starting Text-Based HP SMH](#)” (page 30).
2. Press **u** to select Accounts for Users and Groups.
3. Press **I** to select Local Users or press **n** to select NIS Users.
4. Press **a** to select Add User and fill in the form.
5. **Login Name**

Enter the user’s login name. This must start with a letter. The maximum length can be eight or 255 characters, depending on whether long user names are set. See “[Setting Long User and Group Names](#)” (page 43).

(passwd field 1; shadow field 1)

6. **User ID**

Select a numeric user ID.

If you mark `Next Available ID`, HP SMH will select the next available user ID after 100 (not the next ID after the current highest ID).

If you mark `Specify ID`,

- a. A space is displayed for you to type in the number.
- b. Normally, IDs are expected to be unique (the value for `Allow Duplicate User ID` is `No`).

If you want a duplicate user ID, change the value for `Allow Duplicate User ID` to `Yes`.

(passwd field 3)

7. **Primary Group**

Select the user’s primary group. This defaults to `users`, which is customary for the normal users on a system. You can choose another defined group by typing its name in the space or by selecting **Change Primary Group**, which displays a list of the defined groups to choose from.

To create a group, see “[Adding a Group with Text-Based HP SMH](#)” (page 59).

(passwd field 4)

8. **Home Directory**

Select the user’s home directory. Normally, this is `/home/loginname`, which is selected by the keyword `default`. To change it, enter the full path name of the chosen directory.

(passwd field 6)

9. Create Home Directory

If you want HP SMH to create the home directory, set **Create Home Directory** to Yes. If not, set it to No.

10. Start-Up Program Options

Choose a login shell.

If you check **Select Start-Up Program**, the **Start-Up Program** field offers a list of shells to choose from. By default, if the file /etc/shells does not exist, HP SMH sets the list of shells to:

/sbin/sh	POSIX shell (see <i>sh-posix(1)</i>)
/usr/bin/sh	POSIX shell (see <i>sh-posix(1)</i>)
/usr/bin/rsh	restricted POSIX shell (see <i>sh-posix(1)</i>)
/usr/bin/ksh	Korn shell (see <i>ksh(1)</i>)
/usr/bin/rksh	restricted Korn shell (see <i>ksh(1)</i>)
/usr/bin/csh	C shell (see <i>csh(1)</i>)
/usr/bin/keysh	Key shell (see <i>keysh(1)</i>)

if /etc/shells exists, (see *shells(4)*), only the actual executable file names from that file, plus /sbin/sh, are listed. (/sbin/sh must be used by root.)

If you check **Specify Start-Up Program**, the **Start-Up Program** field lets you enter the name of an executable program that will be used as the shell.

(passwd field 7)

11. Comments

Enter comma-separated information in the field.

This information is placed in what has long been known as the *gecos*¹ or *pw_gecos* field of the entry in the /etc/passwd file. The four subfield names (Real Name, Location, Phone, Home Phone) are used by the finger and passwd commands. The Real Name subfield is often used for identification by other system programs, such as lp. The subfields can contain any data you think is pertinent. Due to security issues, Home Phone is rarely used as such any more.

(passwd field 5)

12. Account Aging Options

If in Shadow Password mode (see “Setting Shadow Password Mode” (page 43)), select one of the options. The choices are:

No Restrictions (Normal Behavior) The account has no restrictions.

(shadow fields 7 and 8)

1. gecos stands for General Electric Comprehensive Operating Supervisor, used on early UNIX systems at Bell Laboratories.

Enable Account Aging

Number of Days of Account Inactivity Allowed : -1_____

Enter the number of days the account can go without a login. If the time between logins is exceeded, the account is disabled. The value -1 disables this restriction.

(shadow field 7)

Account Lifetime (mm/dd/yy)

: _____

Enter the expiration date in two-digit month/day/year format. When that date is passed, the account is disabled. If the field is blank, the account will not expire.

(shadow field 8)

13. Password Aging Options

Select one of the options. The choices are:

No Restrictions (Normal Behavior)

The user can change the password at will.
(passwd field 2; shadow fields 4, 5, 6)

Force Password Change at Next Login

The user must change the password at the next login and thereafter can change the password at will.

(passwd field 2; shadow field 3, 4, 5, 6)

Allow Only Super-User To Change Password

Only a superuser can change the account's password. This is *not* recommended.
(passwd field 2; shadow fields 4, 5)

Enable Password Aging

The following fields are displayed. The values in days are rounded up to the nearest multiple of seven.

Max Time Allowed Between Password Changes (7-441 Days) : 7____

Enter a value in the range. If the time expires, the account is disabled.

(passwd field 2; shadow field 5)

Min Time Required Between Password Changes (0-434 Days) : 0____

Enter a value in the range and less than or equal to the Max Time value. The user cannot change the password until this time expires.

(passwd field 2; shadow field 4)

Number of Days to Warn Before Password Expires (0-434 Days): 0

Displayed only in Shadow Password mode (see “Setting Shadow Password Mode” (page 43)). Enter a value in the range and less than or equal to the Max Time value. When this warn limit is reached, a message is displayed every time the user logs in; for example:

Your password will expire in 77 days.

(shadow field 6)

Force Password Changes on Next Login: (X) No
() Yes

If set to Yes, the user must change the password at the next login and thereafter can change the password according to the Max and Min limits above.

(passwd field 2; shadow field 3)

14. (Optional) Select **Preview** to see the commands that will create the account. Press **OK** to continue.
15. Select **Add** to create the account or **Cancel** to quit the process.
16. If the account is enabled, the password dialog is displayed.

Changing password for *loginname*

New password: **password**

Re-enter new password: **password**

Enter a password at the prompts. While the password can be set to null, this is a security breach. It is better to set a password and have the user change it when the user logs in for the first time, for example, by selecting Force Password Change at Next Login.

(passwd field 2; shadow field 2)

17. HP SMH does the following:
 - Creates an entry for the user in the /etc/passwd file (and in the /etc/shadow file, if Shadow Passwords are enabled).
 - Creates the home directory for the user (if requested).
 - Copies all the files (and their permissions) from the “skeleton” directory (if it exists) to the new home directory (if it exists). See “Skeleton Directory” (page 54).
 - Sets the user and group permissions of the home directory and the copied files to the login name and primary group.

18. When the process completes, you are returned to the **Local User** or **NIS User** listing.

Skeleton Directory

The skeleton directory contains files that are copied to a new home directory by HP SMH and the useradd command. The default skeleton directory is /etc/skel. Files can be added and removed. A different directory can be used; see “[Changing the Skeleton Directory](#)” (page 54). The default files in /etc/skel are shown in Table 4-1.

Table 4-1 Default Files in the Skeleton Directory

File Name	Purpose
.cshrc	Start-up file for the C shell, csh.
.exrc	Start-up file for the text editors, ex and vi.
.login	Start-up file for the C shell, csh.
.profile	Start-up file for the POSIX shell, sh and rsh. Start-up file for the Korn shell, ksh and rksh.

Some suggested or recommended files are shown in Table 4-2.

Table 4-2 Suggested Files for the Skeleton Directory

File Name	Purpose
.kshrc	Conventional start-up file for the POSIX shell, sh and rsh, and the Korn shell, ksh and rksh. The ENV environment variable, which is usually defined in .profile, specifies the name of this file.
.forward	This file is used by sendmail to redirect messages. If the user does not receive e-mail on the system, the file can be edited to point to the correct location.
.rhosts	This file can be edited to allow users on other systems to rlogin to this user's account on this system without a password.

Changing the Skeleton Directory

You can designate a different directory for the account skeleton with the useradd -D -k newskel command (see *useradd(1M)*).

This is useful if you modify the skeleton files or add other files to provide the initial user environment. You can also create different skeletons for different user groups.

By not modifying /etc/skel itself, you retain the original installed information.

Using a Template to Add a User with Text-Based HP SMH

Use this procedure to add a new user with the assistance of an HP SMH user template. If you need to define a template, go to the procedure at “[Making User Templates with Text-Based HP SMH](#)” (page 55), then return here.

1. Start HP SMH, as described in “[Starting Text-Based HP SMH](#)” (page 30).
2. Press **u** to select Accounts for Users and Groups.
3. Press **t** to select Templates. (If there is no template, go to “[Making User Templates with Text-Based HP SMH](#)” (page 55)).
4. Highlight a template name and press **s** to select it.
5. Press **Esc** to return to the previous menu.
6. Press **I** to select Local Users or **n** to select NIS Users.
7. Press **a** to select Add User. The fields that are not predefined by the template are displayed.
8. Enter the user’s **Login Name**.

This must start with a letter and be up to either 8 or 254 characters long, depending on whether long user and group names is set. See “[Setting Long User and Group Names](#)” (page 43).

9. If the **User ID** field is displayed, enter a numeric user ID.
 - a. A space is displayed for you to type in the number.
 - b. Normally, IDs are unique (the value for **Allow Duplicate UID** is **No**). If you want a duplicate user ID, select **Yes**.

If the field is not displayed, HP SMH selects the next available user ID after 100 (not after the highest current ID).

10. Enter comma-separated information in the **Comments** field. See Step 11 in “[Adding a User with Text-Based HP SMH](#)” for details.
11. (Optional) Select **Preview** to see the commands that will create the account. Press **OK** to continue.
12. Select **Add** to create the account or **Cancel** to quit the process.
13. If the template requires a password, enter a password in the password dialog. See Step 16 in “[Adding a User with Text-Based HP SMH](#)” for details.
14. HP SMH creates the user account. See Step 17 in “[Adding a User with Text-Based HP SMH](#)” for details.
15. When the process completes, you are returned to the **Local Users** or **NIS Users** listing.

Making User Templates with Text-Based HP SMH

A template is a way to predefine the contents of most of the fields of a user account so many user accounts with the same parameters can be created with the fewest steps.

These templates are available to both web-based and text-based HP SMH. They can be made with either version. The following instructions described the text-based process.

1. Start HP SMH, as described in “Starting Text-Based HP SMH” (page 30).
2. Press **u** to select **Accounts for Users and Groups**.
3. Press **t** to select **Templates**.
4. Press **a** to select **Add User Template**.
5. At **Template Name**, enter a name for the template of up to 16 characters. This is displayed on the template menu and when you add a user.
6. At **Template Description**, enter a description of the template of up to 50 characters. This is displayed on the template menu and when you add a user.
7. At **UID Generation Method**, select the user ID selection method. If you choose **First Available**, the first available number after 100 is automatically assigned. If you choose **Prompt for it**, the field for the number and the **Allow Duplicates** question will be displayed when you add a user.
8. At **Primary Group Name**, enter a primary group name. The default is **users**.
If you tab to and select the **Change Primary Group** button, the **Select Primary Group** screen is displayed with a list of the current group names. Highlight the one you want and press **s (Select and Go Back)**. You return to the **Add Template** screen with the group value filled in.
The name you choose, whether typed in or selected must exist as a group name when the template is used to add a user. Otherwise, the add will fail.
9. At **Home Directory**, enter a full path name for the parent of the home directory. The user's home directory will be *thisvalue/username*. The default is **/home**.
10. At **Create Home Directory**, select **Yes** or **No** to create the home directory.
11. At **Start-Up Program Options**, choose one of **Select Start-Up Program** or **Specify Start-Up Program**.
If you choose **Select Start-Up Program**, choose a login shell from the drop-down list.
If you choose **Specify Start-Up Program**, enter the login shell in the space provided.
12. At **Comment Settings**, choose a comment setting. If you choose **None**, the comment field will be empty in the */etc/passwd* entry. If you choose **Prompt For It**, the field will be prompted when you add a user.
13. At **Account Status**, choose whether the account will initially be enabled or disabled.
14. At **Account Password**, choose whether the account password will initially be null or will be prompted for when you add a user.

15. At **Account Aging Options**, make the selections as described in Step 12 of “Adding a User with Text-Based HP SMH”.

This information is stored in the template but is used only if Shadow Password mode is set when the user is added.

16. At **Password Aging Options**, make the selections as described in Step 13 of “Adding a User with Text-Based HP SMH”.
17. At **Security Options**, select one of the following:

Use System-Wide Values for Security Attributes

The system-wide security attributes will be applied to the new account. See “Configuring System Default Security Attributes” (page 63). The attributes are also described in *security(4)*.

Specify Security Values

You can provide individual exceptions to the system-wide values for the following attributes. The system-wide default values are displayed. See “Configuring User Security Attributes” (page 65) for details. The attributes are also described in *userdb(4)*.

ALLOW_NULL_PASSWORD	(0 or 1)	:	0 _____
AUDIT_FLAG	(0 or 1)	:	0 _____
AUTH_MAXTRIES	(0-999)	:	0 _____
DISPLAY_LAST_LOGIN	(0 or 1)	:	0 _____
LOGIN_TIMES	(Any)	:	Any _____
MIN_PASSWORD_LENGTH	(6-8)	:	6 _____
NUMBER_LOGINS_ALLOWED	(0-999)	:	0 _____
PASSWORD_HISTORY_DEPTH	(1-24)	:	1 _____
PASSWORD_MIN_LOWER_CASE_CHARS	(0-7)	:	0 _____
PASSWORD_MIN_UPPER_CASE_CHARS	(0-7)	:	0 _____
PASSWORD_MIN_SPECIAL_CHARS	(0-6)	:	0 _____
PASSWORD_MIN_DIGIT_CHARS	(0-6)	:	0 _____
UMASK	(0-511 leading zero denotes octal)	:	0 _____



NOTE: The upper limit for UMASK is shown here in decimal (decimal 511 = octal 0777). A leading zero is necessary to specify octal here.

Modifying a User with Text-Based HP SMH

1. Start HP SMH, as described in “Starting Text-Based HP SMH” (page 30).
2. Press **u** to select **Accounts for Users and Groups**.
3. Press **1** to select **Local Users** or **n** to select **NIS Users**.
4. Highlight the login name you want to modify and press **m**.

5. You can modify the following data by typing in new values or making different selections. See “Adding a User with Text-Based HP SMH” (page 50) for details on the fields and selections.
 - **Login Name**
 - **User ID**
 - **Allow Duplicate User ID**
 - **Primary Group**
 - **Home Directory**
 - **Create Home Directory**
 - **Login Shell**
 - **Comments**
 - **Account Options**
 - **Password Options**
6. (Optional) Select **Preview** to see the commands that will modify the account. Press **OK** to continue.
7. Select **Modify** to change the user or **Cancel** to quit the process.
8. HP SMH does the following:
 - Makes appropriate changes in the entry for the user in /etc/passwd.
 - Creates the new home directory for the user, if the **Home Directory** name is altered.
 - Copies the contents of the old home directory to the new home directory, if the **Home Directory** name is altered and **Create Home Directory** is set to **Yes**.

The old home directory and its files remain unchanged.
 - Sets the user and group ownership of the home directory and the copied files to the login name and primary group, as necessary.
 - Changes the user ID of all the user's files throughout the system, if the **User ID** is changed.
9. When the process completes, you are returned to the **Local User** or **NIS User** listing.

Deleting a User with Text-Based HP SMH

1. Start HP SMH, as described in “Starting Text-Based HP SMH” (page 30).
2. Press **u** to select **Accounts for Users and Groups**.
3. Press **1** to select **Local Users** or **n** to select **NIS Users**.
4. Highlight the login name you want to delete and press **r**.

HP SMH displays a screen that asks what to do with the user's files and directories. Select one of the following choices:

Leave Files Undisturbed	None of the files or directories owned by the user on the system will be modified, except that listings will show the user ID, not the user name.
Remove from User's Home Directory Only	All the files owned by the user beneath the user's home directory will be deleted. The home directory and subdirectories owned by the user will be deleted if they are empty.
Remove from All Local File Systems	None of the files or directories owned by the user elsewhere on the system will be modified, except that listings will show the user ID, not the user name.
Reassign to a Specified User	All files owned by the user will be deleted. All directories owned by the user, including the home directory, will be deleted if they are empty.

5. (Optional) Select **Preview** to see the commands that will remove the account. Press **OK** to continue.
6. Select **Delete** to delete the user or **Cancel** to quit the process.
7. HP SMH removes the account entry from /etc/passwd and deletes or changes ownership of files and directories as described above.
8. When the process completes, you are returned to the **Local User** or **NIS User** listing.

Adding a Group with Text-Based HP SMH

1. Start HP SMH, as described in “Starting Text-Based HP SMH” (page 30).
2. Press **u** to select **Accounts for Users and Groups**.
3. Press **g** to select **View or Configure Groups**.

The current list of groups is displayed with columns for the group name, the group ID, and the user names that have the group as a secondary group.

4. Press **a, Add Group**, and fill in the blanks.

5. **Group Name**

Enter the group name. This must start with a letter. The maximum length can be 16 or 255 characters, depending on whether long group names are set. See “Setting Long User and Group Names” (page 43).

(group field 1)

6. **Group ID**

Select a numeric group ID.

If you mark **Next Available ID**, HP SMH will select the next available user ID after 100 (not the next ID after the current highest ID).

If you mark **Specify ID**,

- a. A space is displayed for you to type in the number.
- b. Normally, IDs are expected to be unique (the value for **Allow Duplicate User ID** is **No**).

If you want a duplicate user ID, change the value for **Allow Duplicate User ID** to **Yes**.

(group field 3)

7. **Users with this Group as Secondary Group**

Scroll through the list of user names and mark those that you want to have this group as a secondary group.

8. (Optional) Select **Preview** to see the commands that will add the group. Press **OK** to continue.
9. Select **Add** to add the group or **Cancel** to quit the process.

Managing Users and Groups with Commands

Automating the Process of Adding a User

When you have several users to add to a system, you can save time by:

- Using the HP SMH Template; see “Using a Template to Add a User with Text-Based HP SMH” (page 55).
- Using the `useradd` Command; see “Using the `useradd`, `usermod`, and `userdel` Commands” (page 60).

Using the `useradd`, `usermod`, and `userdel` Commands

You can use the `useradd` command to add users, `usermod` to modify them, and `userdel` to delete them. See the *useradd(1M)*, *usermod(1M)*, and *userdel(1M)* manpages.

`useradd` has the form:

```
/usr/sbin/useradd [option]... username
```

username is the login name for the new user.

Some of the options are described in Table 4-3. For all the options and complete information on the command, see *useradd(1M)*.

Table 4-3 *useradd Options*

Option	Meaning
<code>-b <i>b_dir</i></code>	Default base directory for user home directory. The default is <code>/home</code> .
<code>-c "comments"</code>	Full name or other comments. This is often a comma-separated string in the form: <i>fullname, location, workphone, homephone</i>
<code>-d <i>dir</i></code>	Home directory path name. The default is <code><i>b_dir/username</i></code> .
<code>-e <i>date</i></code>	Account expiration date. The default is none. To use the <code>-e</code> option, you must enable shadow passwords. For details on how to do that, see <i>pwconv(1M)</i> .
<code>-f <i>n</i></code>	Number of days the account can be inactive before being disabled. As with the <code>-e</code> option, to use the <code>-f</code> option you must enable shadow passwords. For details on how to do that, see <i>pwconv(1M)</i> .
<code>-g <i>group</i></code>	Primary working group name or group ID. Group must exist. The default is <code>users</code> (group ID 20).
<code>-G <i>groups</i></code>	Comma-separated list of secondary groups. Groups must exist.
<code>-k <i>skeldir</i></code>	Skeleton directory containing initialization files. The default is <code>/etc/skel</code> .
<code>-m</code>	Create the home directory in addition to defining user. The default is don't create home directory.
<code>-s <i>shell</i></code>	Shell. The default is <code>/sbin/sh</code> .
<code>-u <i>uid</i></code>	User ID. The default is the first available number after 100.

The following command creates a new user account, adds `patrick` to the primary working group (called `users`), creates a home directory, and sets up a default Korn shell:

```
# useradd -g users -m -k /etc/skel -s /usr/bin/ksh patrick
```

The resulting entry in the `/etc/passwd` file is:

```
patrick:*:104:20::/home/patrick:/usr/bin/ksh
```

You can make a script with as many instances of the `useradd` command as necessary. You can set different defaults with the `useradd -D` command.

After the accounts are created, set their initial passwords with the `passwd` command.

Manually Adding a User

Use the following steps to add a user from the command line.

1. Add the user to the /etc/passwd file.

As root, use the /usr/sbin/vipw command to edit /etc/passwd. See *vipw(1M)*, *passwd(4)*, and *passwd(1)*.

For example, you might want to add this line for user tom:

```
tom:*:102:20:Tom,,,:/home/tom:/usr/bin/sh
```

This creates the entry and disables logins (the * in the password field). The home directory is /home/tom and the login shell is /usr/bin/sh. The user ID is 102 and the primary group ID is 20, conventionally, users.

2. Use the passwd command to set an initial password for the account. For example:

```
# passwd tom
Changing password for tom
New password: password
Re-enter new password: password
passwd successfully changed
```

3. Use the passwd command to force a password change at the next login. For example:

```
# passwd -f tom
```

4. Create a home directory. For example:

```
# /usr/bin/mkdir /home/tom
```

Change the ownership of the directory to the user's name. For example:

```
# /usr/bin/chown tom:users /home/tom
```

5. Ensure that the user has the appropriate shell start-up files to execute when logging in.

You can create standard start-up files (templates) that can be copied to users' directories. The directory most often used for this purpose is /etc/skel. See "Skeleton Directory" (page 54).

For example:

```
# cp /etc/skel/.profile /users/tom/.profile
```

6. Change the ownership of the start-up file to the new user's account and group. For example:

```
# /usr/bin/chown tom:users .profile
```

Configuring System and User Security

Configuring System Default Security Attributes

1. Start HP SMH:
 - **Using the web-based version of the HP SMH:**
 - a. Use the URL: `http://your_system:2301` to start the web-based interface in your browser.
 - b. Log in using the user name and password of an account with the appropriate privileges (usually root)
 - **Using the text-based interface:**
 - a. Enter the command: `/usr/sbin/secweb -t` to start the text-based interface.
2. Navigate to the System Defaults page:
 - **Using the web-based version of the HP SMH:**
Select **Tools → Auditing and Security Attributes Configuration(web-based) → System Defaults**
 - **Using the text-based interface:**
Press **c** to select **Security Attributes Configuration**. Then press **s** to select **System Defaults**.
3. In the text-based version of the interface, the table in Figure 4-1 is displayed. It shows each attribute's name, its default value, and its current setting. The individual attributes are described in the *security(4)* manpage.
In the web-based version of the interface, a similar list is displayed.

Figure 4-1 Security Attributes Configuration: System Defaults

Attribute	Default	Value
ABORT_LOGIN_ON_MISSING_HOMEDIR	0	<default>
ALLOW_NULL_PASSWORD	1	<default>
AUDIT_FLAG	1	<default>
AUTH_MAXTRIES	0	<default>
BOOT_AUTH	0	<default>
BOOT_USERS	root	<default>
DISPLAY_LAST_LOGIN	1	<default>
INACTIVITY_MAXDAYS	0	<default>
LOGIN_TIMES	Any	<default>
MIN_PASSWORD_LENGTH	6	<default>
NOLOGIN	0	<default>
NUMBER_OF_LOGINS_ALLOWED	0	<default>
PASSWORD_HISTORY_DEPTH	1	<default>
PASSWORD_MIN_LOWER_CASE_CHARS	0	<default>
PASSWORD_MIN_UPPER_CASE_CHARS	0	<default>
PASSWORD_MIN_DIGIT_CHARS	0	<default>
PASSWORD_MIN_SPECIAL_CHARS	0	<default>
PASSWORD_MAXDAYS	-1	<default>
PASSWORD_MINDAYS	0	<default>
PASSWORD_WARMDAYS	0	<default>
SU_DEFAULT_PATH	{null}	<default>
SU_KEEP_ENV_VARS	{null}	<default>
SU_ROOT_GROUP	{null}	<default>
UMASK	0	<default>

4. To view more information about an attribute:

- **Using the web-based version of the HP SMH:**

Click on the attribute you want information about. Details will be displayed at the bottom of your browser window.

- **Using the text-based interface:**

Highlight the attribute and press **Enter**. For example, for the NOLOGIN attribute, the screen would show:

```

Attribute      NOLOGIN
Description   Can /etc/nologin be used to disable non-root logins? (0=No 1=Yes)
Min Value    0
Max Value    1
Default      0
Value        0

```

5. To *modify* the value:

- **Using the web-based version of the HP SMH:**

With the desired attribute highlighted (you clicked on it), select » **Modify System Value ...** on the right-hand side of the display.

A new page will be displayed with the description and current value for the attribute.

Enter a new value for the attribute and click on the **Modify** button at the bottom of the display.



NOTE: To preview what command will be executed by HP SMH before you click on **Modify**, you can click on **Preview**.

- **Using the text-based interface:**

press **m**. For NOLOGIN, the screen would show (slightly condensed):

Modify the system value by entering a valid value as specified in security(4) man page. Enter default to reset the system value to the default value.

Note: The HP-UX Security Attributes Configuration Tool only checks for valid ranges. It does not perform any checks to ensure the correctness of the value entered.

```
Attribute      : NOLOGIN
Description    : Can /etc/nologin be used to disable non-root logins? (0=No 1=Yes)
Range         : 0...1
Default       : 0
System Value   : 0 _____
[ Modify ] [ Preview ] [ Cancel ] [ Help ]
```

Enter an appropriate value for **System Value**. To choose the default value, enter **default**.

(Optional) Select **Preview** to see the commands that will change the value. Press **OK** to continue.

Select **Modify** to change the value or **Cancel** to quit the process.

Configuring User Security Attributes



NOTE: You can also access this procedure from the **Accounts for Users and Groups** function. On the **Local Users** or **NIS Users** screen, highlight an account and press **s**, **Modify Security Attributes**. Then continue below with Step 4.

1. Start HP SMH, as described in “Starting Text-Based HP SMH” (page 30).
2. Press **c** to select **Security Attributes Configuration**.
3. Press **1** to select **Local Users** or **n** to select **NIS Users**. A list of users is displayed. The **User Values** column indicates whether any user values have been specified for the user.

```

Displaying Local Users
Name          User Id    User Values
=====
adm           4          no
allamp        1834       no
anewuser      111         yes
bin           2          no

```

4. Highlight a user and press **Enter**.

The table in Figure 4-2 is displayed. It shows the attribute name, the current setting for the user (- means the system value), and the current system value. System defaults are marked with the word Default and the default value in parentheses, as in Default (1). The individual attributes are described in the *security(4)* manpage.

Figure 4-2 Security Attributes Configuration: Local or NIS Users

Attribute	User Value	System Value
ALLOW_NULL_PASSWORD	-	Default (1)
AUDIT_FLAG	-	0
AUTH_MAXTRIES	-	Default (0)
DISPLAY_LAST_LOGIN	-	Default (1)
INACTIVITY_MAXDAYS	-	Default (0)
LOGIN_TIMES	-	Default (Any)
MIN_PASSWORD_LENGTH	-	Default (6)
NUMBER_OF_LOGINS_ALLOWED	-	Default (0)
PASSWORD_HISTORY_DEPTH	-	Default (1)
PASSWORD_MIN_LOWER_CASE_CHARS	-	Default (0)
PASSWORD_MIN_UPPER_CASE_CHARS	-	Default (0)
PASSWORD_MIN_DIGIT_CHARS	-	Default (0)
PASSWORD_MIN_SPECIAL_CHARS	-	Default (0)
PASSWORD_MAXDAYS	-	Default (-1)
PASSWORD_MINDAYS	-	Default (0)
PASSWORD_WARNDAYS	-	45
UMASK	-	Default (0)



NOTE: INACTIVITY_MAXDAYS and PASSWORD_WARNDAYS are only displayed if the system is in Shadow Password mode.

5. To see more detail for an attribute, highlight the attribute and press **Enter**. The additional data includes a description and the minimum and maximum values. Press **Esc** to return to the attribute list.
6. To set or change the values for the user, press **c**, **Configure Per User Exceptions**. The **Configuration** screen is displayed. It includes all the attributes. The following is an abbreviated view:

```

-----
Attribute Name          [Range, System Value]
-----
```

```
(Description)
-----
ALLOW_NULL_PASSWORD          [0...1 , 1 ]: default _____
(Allow login with null password? (0=No 1=Yes))
...
UMASK                      [0...511, 0 ]: default _____
(Default umask (leading zero denotes octal value))

[ Modify ] [ Preview ] [ Cancel ] [ Help ]
```

7. Make the changes you want in the spaces provided.
8. (Optional) Select **Preview** to see the commands that will modify the attributes. Press **OK** to continue.
9. Select **Modify** to change the attributes or **Cancel** to quit the process.
10. The changes from the default values are shown in the **User Value** column.

Controlling File Access

Working groups, file permissions, and file ownership all determine who can access a given file. See also the *HP-UX System Administrator's Guide: Security Management*.

Defining Group Membership

Users on your system can be divided into working groups so that files owned by members of a given group can be shared and yet remain protected from access by users who are not members of the group. A user's primary group membership number is included as one entry in the /etc/passwd file. Group information is defined in /etc/group and /etc/logingroup.

Users who are members of more than one group, as specified in /etc/group, can change their current group with the /usr/bin/newgrp command. You do not need to use the newgrp command if user groups are defined in /etc/logingroup. If you do not divide the users of your system into separate working groups, it is customary to set up one group (usually called *users*) and assign all users of your system to that group.

You can use HP SMH to add, remove, or modify group membership.

To manually change group membership, edit /etc/group and optionally /etc/logingroup with a text editor, such as vi. Although you can enter a group-level password in /etc/group, it is not recommended. To avoid maintaining multiple files, you can link /etc/logingroup to /etc/group. For details on the /etc/group and /etc/logingroup files, see the *group(4)* manpage. For information on linking files, see the *link(1M)* and *ln(1)* manpages.

You can assign special privileges to a group of users using the /usr/sbin/setprivgrp command. For more information, see *chown(1)*, *getprivgrp(1)*, *setprivgrp(1M)*, *chown(2)*, *getprivgrp(2)*, *lockf(2)*, *plock(2)*, *rtprio(2)*, *setgid(2)*, *setgid(2)*, *setprivgrp(2)*, *setuid(2)*, *shmctl(2)*, and *shmctl(2)*.

Setting File Access Permissions

The `/usr/bin/chmod` command changes the type of access (read, write, and execute privileges) for the file's owner, group members, or all others. Only the owner of a file (or the superuser) can change its read, write, and execute privileges. For details, see *chmod(1)*.

By default, new files have read/write permission for everyone (`-rw-rw-rw-`) and new directories have read/write/execute permission for everyone (`drwxrwxrwx`). Default file permissions can be changed using the `/usr/bin/umask` command. For details, see *umask(1)*. The default for trusted systems is different; see the *HP-UX System Administrator's Guide: Security Management*.

Setting Ownership for Files

The `/usr/bin/chown` command changes file user (and group) ownership. To change the user, you must own the file (and belong to a group with the CHOWN privilege) or have superuser privileges.

The `/usr/bin/chgrp` command changes file group ownership. To change the group, you must own the file (and belong to a group with the CHOWN privilege) or have superuser privileges.

For more information, refer to *chown(1)* and *chgrp(1)*.

Setting Access Control Lists

Access control lists (ACLs) offer a finer degree of file protection than traditional file access permissions. You can use ACLs to allow or restrict file access to individual users unrelated to what group the users belong. Only the owner of a file (or the superuser) can create ACLs.

ACLs are supported on both JFS and HFS file systems, but the commands and some of the semantics differ. On a JFS file system, use `setacl` to set ACLs and use `getacl` to view them. On an HFS file system, use `chacl` to set ACLs and use `lsacl` to view them.

For a discussion of both JFS and HFS ACLs, see the *HP-UX System Administrator's Guide: Security Management*.

For additional JFS ACL information see *setacl(1)*, *getacl(1)*, and *aclv(5)*.

For additional HFS ACL information, see *lsacl(1)*, *chacl(1)*, and *acl(5)*.

Customizing System-Wide and User Login Environments

Defaults for system-wide variables, such as time-zone setting, terminal type, search path, and mail and news notification, can be set in `/etc/profile` for Korn and POSIX shell users and in `/etc/csh.login` for C shell users.

User login scripts can be used to override the system defaults. When HP SMH adds a user, default user login scripts are copied to the user's home directory. For Korn and

POSIX shell users /etc/skel/.profile is copied to the home directory as .profile. For C shell users, /etc/skel/.login and /etc/skel/.cshrc are copied to the home directory as .login and .cshrc. Refer to the *Shells: User's Guide* and the *Technical Addendum to the Shells: User's Guide* for information on customizing user login scripts.



NOTE: Do a full backup once you have initially set up and customized your system. This allows you to reconstruct your system — kernel, system files, file system structure, user structures, and your customized files — if you need to. Use HP-UX commands to perform the backup, as described in the *HP-UX System Administrator's Guide: Routine Tasks*.

Accessing Multiple Systems

If a user has an account with the same login on more than one system (for example, if the user's home directory is NFS-mounted from a file server), the user ID number should be the same on all of these systems.

For example, suppose user tom has a user ID of 200 on system dept27 and shares files to wsj6700 where he has a user ID of 330. If the files created on dept27 have permissions of -rw-----, then they will not be accessible to him from wsj6700. HP-UX determines file ownership by the user ID, not by the user name.

As system administrator, you need to ensure that each new user login name has a corresponding user ID that is unique within the workgroup, site, or network that the user needs to reach.

For information on whether you should share users' home and mail directories, see the *HP-UX System Administrator's Guide: Overview*.

To allow a user to access a remote system with rcp, remsh, or rlogin without supplying a password, set up \$HOME/.rhostsfile on the remote system. See “\$HOME/.rhosts File” (page 72).

Consider using the Network Information Service (NIS) to manage your users on multiple systems. See the *NIS Administrator's Guide*.

Sharing Remote Work Directories

After you have created a new user's account, you must decide which directories within the workgroup the user needs to access. NFS allows users to use their own local systems to work on files residing on file servers or other systems in the workgroup. The server or remote system shares with the local system and the local system mounts from the remote system.

The topic “Adding a User to Several Systems: A Case Study” (page 70) illustrates how you might set up your users.

Local versus Remote Home Directories

Users can have their home directory on their own local system or on a remote file server. The advantage of keeping all users' home directories on one file server is that you can back up all the accounts at one time.


If a user's home directory is on a remote server, you may want to create a minimal home directory on the local system so that a user can still log into the local system if the server is down. For information on whether you should share users' home and mail directories, see the *HP-UX System Administrator's Guide: Overview*.

See "Adding a User to Several Systems: A Case Study" (page 70) for steps to create a home directory on a remote system.

Adding a User to Several Systems: A Case Study

The following example shows how to mount Tom's home directory and work directory from the file server, flserver, and mount Emacs from the application server, appserver.

Figure 4-3 Adding a User to Several Systems



Before beginning, make sure Tom's login name has a user ID that is unique across the systems he is going to use. (Your network administrator may have a program to ensure the uniqueness of user ID numbers.)

Next, create an account for Tom on the file server, `f1server`. See “Configuring Users and Groups with HP SMH” (page 49).

Then, perform the following procedure:

1. On the file server, share Tom’s home directory and the directory where he does his work:
 - a. Add an entry in the `/etc/dfs/dfstab` file to share Tom’s home directory:

```
share -F nfs -o access=appservr:dept27:wsj6700 /home/tom
```

If the directory is already shared, simply add Tom’s system to the access list.
 - b. Add an entry to the `/etc/dfs/dfstab` file to share the `/work` directory:

```
share -F nfs -o access=dept27:wsj6700 /work
```

This contains the files and directories Tom will share with other members of his project team.
 - c. Force the server to re-read `/etc/dfs/dfstab` and activate the new shares for `/work` and `/home`:
`# /usr/sbin/shareall`
2. On the application server, share the Emacs directories that Tom needs:
 - a. Add entries to the `/etc/dfs/dfstab` file:

```
share -F nfs -o access=dept27:wsj6700 /usr/local/share/emacs  
share -F nfs -o access=dept27:wsj6700 /opt/hp/gnu/bin700/emacs
```
 - b. Share the directories for Emacs:
`# /usr/sbin/shareall`
3. On Tom’s login server, `dept27`, do the following:
 - a. Create Tom’s account. See “Configuring Users and Groups with HP SMH” (page 49). If Tom’s login has already been set up on another system (for example on `f1server`), you may want to cut the line from `f1server`’s `/etc/passwd` file and paste it into the `/etc/passwd` file on `dept27` to ensure that Tom’s account has the same user name and user ID on both systems.
 - b. Create empty directories for the file systems to be mounted.
`# mkdir /home/tom
mkdir /work
mkdir /usr/local/share/emacs
mkdir /opt/hp/gnu/bin700/emacs`
 - c. Add entries to `/etc/fstab`.
`f1server:/home/tom /home/tom nfs rw,suid 0 0
f1server:/work /work nfs rw,suid 0 0
appserver:/usr/share/emacs /usr/share/emacs nfs rw,suid 0 0
appserver:/opt/hp/gnu/bin700/emacs nfs rw,suid 0 0`
 - d. Mount all the directories:
`# mount -a`

See “Sharing an HP-UX Directory” (page 77) for more information.

Sharing a Local Home Directory

Assume you are setting up an account on the system named wsj6700 for the user lisa. In this example, lisa’s home directory will reside on her local disk and will be shared with the other systems she logs in to.

1. On the local system, do the following:
 - a. Create the user’s account. See “Configuring Users and Groups with HP SMH” (page 49).
 - b. Share the user’s home directory with other systems that the user needs to log in to:
 - Add an entry, such as flserver, to /etc/dfs/dfstab:

```
share -F nfs -o access=mailserver:appserver:flserver /home/lisa
      • Share the home directory/home/lisa:
      # /usr/sbin/share /home/lisa
```

2. On the remote system, do the following:

- a. Create an empty directory:
mkdir /home/lisa
- b. Add entry to /etc/fstab:
wsj6700:/home/lisa /home/lisa nfs rw,suid 0 0
- c. Mount all directories:
mount -a

See “Sharing an HP-UX Directory” (page 77) for more information.

Allowing Access from Remote Systems

To allow a user access from a remote system using rcp, remsh, or rlogin without supplying a password, set up an /etc/hosts.equiv or \$HOME/.rhosts file on the local system. See the *hosts.equiv(4)* manpage for more information.

The /etc/hosts.equiv file can contain NFS netgroups. See the *NFS Services Administrator’s Guide* for more information.

\$HOME/.rhosts File

Users listed in \$HOME/.rhosts are allowed access to the local system, from the remote systems and accounts named in the file, without supplying a password. This file should be owned by the local user.

In the following example, /users/spence/.rhosts resides on system wsj6700. Users tom and patrick can log in to spence's account on wsj6700, from dept27 and wsb2600 respectively, without supplying a password.

```
dept27 tom  
wsb2600 patrick
```

Enabling X Window Server Access

To allow an X Window client to send output to an X Window server using the display option, use the xhost command.

For example, to allow system dept27 to send a window to system wszx6, enter:

```
xhost +dept27
```

on system wszx6. See *xhost(1)* for further details.

5 Configuring Networking

This chapter describes the following networking topics:

- “Configuring the Network File System (NFS)” (page 75)
- “Configuring File Transfer Protocol (FTP)” (page 87)
- “Interfacing with Microsoft Windows Systems” (page 106)

Other networking issues are described in:

- “Setting System and Network Parameters” (page 41)
- “Ethernet Configuration and Verification” (page 135)

Configuring the Network File System (NFS)

This section provides procedures and troubleshooting information for the Network File System (NFS).

NFS allows a computer to access a file system that resides on another computer’s disks, as though the file system were installed locally.

The **NFS server** is the computer to which the disk is physically attached. **NFS clients** are the computers that use the file system remotely. Before an NFS client can **mount** a file system that resides on the NFS server’s disks, the NFS server must **share** it.

Before you can share file systems, you must install and configure NFS software on both the server and client systems. In most cases this will have been done when the systems were installed. Use the *NFS Services Administrator’s Guide* if you need to install NFS.

For information and guidelines on planning a workgroup’s file-sharing configuration, see the *HP-UX System Administrator’s Guide: Overview*.

This section contains information on the following:

- “Sharing an HP-UX Directory” (page 77)
- “Mounting a Shared File System (HP-UX to HP-UX)” (page 78)
- “Troubleshooting NFS” (page 84)
- “Recovering Network Services after a Power Failure” (page 86)
- “Moving or Reusing a Shared Directory” (page 87)

See also:

- “Adding a User to Several Systems: A Case Study” (page 70)

Exporting versus Sharing

Prior to HP-UX 11i v3, file systems were “exported” for use by other systems, using the `exportfs` command. Exported file system information was stored in the `/etc/exports` file.

Beginning with HP-UX 11i v3, file systems are “shared” with other systems with the share command. Shared file system information is stored in the /etc/dfs/dfstab file. For conversion information, see the *HP-UX 11i Version 3 Release Notes*.

Enabling and Disabling the NFS Server and Client

The following procedures describe how to enable or disable the NFS server and client.

Enable or Disable the NFS Server with Text-Based HP SMH

1. Log in to the *server* as superuser.
2. Start HP SMH; see “Starting Text-Based HP SMH” (page 30).
3. Navigate to **Network Services**.
 - a. Press **n, Networking and Communications**
 - b. Press **s, Network Services Configuration**
 - c. Press **k, Network Services**
4. Highlight **NFS Server**.
5. Press **Tab A** to pull down the **Actions** menu.
6. Select one of:
 - Disable** To stop the NFS server.
 - Enable** To start the NFS server.
 - Restart** To restart the NFS server.

Enable or Disable the NFS Client with Text-Based HP SMH

1. Log in to the *client* as superuser.
2. Start HP SMH; see “Starting Text-Based HP SMH” (page 30).
3. Navigate to **Network Services**.
 - a. Press **n, Networking and Communications**
 - b. Press **s, Network Services Configuration**
 - c. Press **k, Network Services**
4. Highlight **NFS Client**.
5. Press **Tab A** to pull down the **Actions** menu.
6. Select one of:
 - Disable** To stop the NFS client.
 - Enable** To start the NFS client.
 - Restart** To restart the NFS client.

Sharing an HP-UX Directory

Use either of the following procedures to set up NFS shares on the server.

- “Share a Directory Using Text-Based HP SMH” (page 77).
- “Share a Directory Using HP-UX Commands” (page 77).



NOTE: An NFS server can share an ordinary file as well as a directory. In either case, the NFS client must mount the shared file system on a directory.

Share a Directory Using Text-Based HP SMH

1. Log in to the *server* as superuser.
2. Start HP SMH, as described in “Starting Text-Based HP SMH” (page 30).
3. Enable the NFS server, if necessary, as described in “Enabling and Disabling the NFS Server and Client” (page 76).
4. Navigate to the **Share/Unshare File System** screen.
 - a. Press **n, Networking and Communications**.
 - b. Press **s, Network Services Configuration**.
 - c. Press **f, Networked File Systems**.
 - d. Press **s, Share/Unshare File System**.

The currently defined shared directories are displayed with columns indicating:

Local Directory	The full path of the local directory name of the file system.
------------------------	---

Currently Shared	Whether the file system is currently shared; is it an entry in /etc/dfs/sharedtab?
-------------------------	--

Permanently Shared	Is it an entry in /etc/dfs/dfstab?
---------------------------	------------------------------------

Logging Enabled	Is the use of the file system being logged.
------------------------	---

5. Press **s, Share a File System**.

A screen showing the needed fields is displayed.

Enter values as described in the *share(1M)* and *share_nfs(1M)* manpages.

6. (Optional) Select **Preview** to see the commands that will be executed.
7. Select **OK** to share the directory or **Cancel** to quit the process.

Share a Directory Using HP-UX Commands

1. Log into the *server* as superuser.

2. If the system is not already configured as an NFS server:
 - a. Edit `/etc/rc.config.d/nfsconf`, setting the following values:

```
NFS_CORE=1  
NFS_SERVER=1  
START_MOUNTD=1
```



NOTE: You can also use the `setoncenv` command to set and display NFS and other Open Network Computing configuration variables. See *setoncenv(1M)*.

- b. Run the `nfs.server` script:

```
# /sbin/init.d/nfs.server start
```

3. Edit `/etc/dfs/dfstab`, adding an entry for each directory that is to be shared. The entry identifies the directory and (optionally) the systems that can import it. The entry should look something like this:

```
share -F nfs -o access=dept27:wsj6700 /opt/hp/gnu/bin700/emacs
```

See *share_nfs(1M)* for more information on the `-o` suboptions `access`, `ro`, and `rw`.

4. To share all the entries in `/etc/dfs/dfstab`:

```
# /usr/sbin/shareall
```

Or to just share the new file system:

```
# /usr/sbin/share /opt/hp/gnu/bin700/emacs
```

Mounting a Shared File System (HP-UX to HP-UX)

Before you begin, you need to:

- Check that the directory on the local (client) system that you are mounting *on* either:
 - Does not already exist; *or*
 - Is empty; *or*
 - Contains data that will not be needed as long as the remote directory is mounted.

In this case, make sure that no one has open files in the local directory and that it is not anyone's current working directory. For example, if you intend to mount on a directory named `/mydir`, enter the following on the client:

```
# fuser -cu /mydir
```



NOTE: Files in the local directory will be hidden, but not overwritten, when you mount the remote directory. The local files will be accessible again once you unmount the remote directory.

- Make sure that the client has permission to share the file system from the server. This requires an entry in `/etc/dfs/dfstab` on the server; see Step 3 under “Share a Directory Using HP-UX Commands” (page 77).
- Decide what type of mount you want. See Table 5-1: “Deciding Which Type of NFS Mount to Use” (page 80).
 - A standard NFS mount. Use one of these procedures:
 - “Standard-Mount a Shared File System Using Text-Based HP SMH” (page 80).
 - “Standard-Mount a Shared File System Using HP-UX Commands” (page 82).
 - An automounted NFS file system using AutoFS. Use this procedure:
 - “Automount a Shared File System Using Text-Based HP SMH” (page 82).

Table 5-1 Deciding Which Type of NFS Mount to Use

You can use a standard NFS mount or the NFS automounter.	
<i>Standard NFS Mount –</i>	Use a standard NFS mount when you would like the mounted file system to always remain mounted. This is useful when the mounted file system will be frequently accessed.
<i>Automounted NFS –</i>	<p>Use an automatically mounted NFS file system when you want the file system to be mounted only when it is actively being used. This is useful when the file system being mounted is used infrequently.</p> <ul style="list-style-type: none"><input type="checkbox"/> AutoFS can be used to mount any type of file system.<input type="checkbox"/> With AutoFS, the configured mount points are the actual mount points.<input type="checkbox"/> You do not have to stop AutoFS to change your automounter maps. The AutoFS daemon, <code>automountd</code>, runs continuously. When you make a change to an automounter map, you run the <code>automount</code> command, which reads the maps, then exits.<input type="checkbox"/> If you use the automounter, the file system will be mounted on the client only when a user or process requests access to it. By default, it will be unmounted after it has remained untouched for five minutes. This default can be changed with the <code>-t</code> option of <code>automount</code> or by setting a value for <code>AUTOMOUNT_TIMEOUT</code> in <code>/etc/default/autofs</code>.<input type="checkbox"/> If you use the automounter <code>-hosts</code> map, HP SMH will create a directory (<code>/net</code> by default) under which all the file systems (on any host on the network) which this client is allowed to import, become available on demand.<input type="checkbox"/> AutoFS map management can be distributed, via a distributed name service, thus not requiring the administrator to modify the <code>/etc/fstab</code> file on every client.

For more information on how to use automounted file systems, see the *NFS Services Administrator's Guide*.



NOTE:

You do not have to call the directory on the client by the same name it has on the server, but it will make things simpler (more transparent) for your users if you do. If you are running applications configured to use specific path names, you *must* make sure those path names are the same on every system on which the applications run.

Standard-Mount a Shared File System Using Text-Based HP SMH



NOTE: The **Disks and File Systems** functional area performs standard mounts. The **Networking and Communications** functional area performs automounts.

1. Log in to the client as superuser.
2. Start HP SMH, as described in "Starting Text-Based HP SMH" (page 30).
3. Enable the NFS client, if necessary, as described in "Enabling and Disabling the NFS Server and Client" (page 76).
4. Navigate to the **Add A New NFS File System** screen.

- a. Press **f, Disks and File Systems**.
- b. Press **f, File Systems - View or Manage File Systems**.
- c. Press **n, Add NFS**.

The **Add A New NFS File System** screen is displayed

5. Fill in the fields. (The default values are marked.)
 - **Mount Point:** The full name of the local directory.
 - **Remote Server:** The full name of the system sharing the file system.
 - **Remote Directory:** The full name of the shared file system.
 - **Mount method:** Check one.
 - Only mount (do not store any config in `/etc/fstab`)
 - Save config in `/etc/fstab` (will not be mounted)
 - Mount now and save config in `/etc/fstab` (default)
 - Check all that apply. (The keywords in parentheses appear in `/etc/fstab`.)
 - Mount read-only (`ro/rw`)
If this is not checked, the file system will be mounted read-write. If it is checked, the file system will be mounted read only.
If the file system is shared read-only, it is a good idea to check this.
 - Do not auto mount (`noauto`)
If this is not checked, the file system is mounted automatically when the system boots. If it is checked, you will have to mount it manually with the `mount` command.
This is *not* related to NFS automounting.
 - Enable SUID (`suid/nosuid`) (default)
If this is checked, programs on the shared file system that have their setuid bit set will run under the program's user ID. If this is not checked, the programs will run under the user's user ID.
 - Enable Quota (`quota/noquota`) (default)
If this is checked, the local system enforces its `quota(1)` quotas. If it is not checked, the local quotas are not enforced. Quotas on the server are always enforced.

The rest of the options are described in the `mount_nfs(1M)` manpage. Usually, they are left unchanged.

6. (Optional) Select **Preview** to see the commands that will mount the file system. Press **OK** to continue.
7. Select **New NFS** to execute the commands or **Cancel** to quit the process.



NOTE: The **Disks and File Systems** functional area performs standard mounts. The **Networking and Communications** functional area performs automounts.

1. Log in to the *client* as superuser.
2. Start HP SMH, as described in “Starting Text-Based HP SMH” (page 30).
3. Enable the NFS client, if necessary, as described in “Enabling and Disabling the NFS Server and Client” (page 76).
4. Navigate to the **Networked File Systems (Automounter)** screen.
 - a. Press **n, Networking and Communications**.
 - b. Press **s, Network Services Configuration**.
 - c. Press **f, Networked File Systems**.
 - d. Press **a, Automounted Remote File Systems**.

The **Networked File Systems (Automounter)** screen is displayed.

5. The currently defined shared directories are displayed with columns indicating:

Mount Directory	The full path of the local directory name of the file system.
Type	Auto.
Remote Server	The full host name of the server.
Remote Directory	The full path of the file system on the remote server.
Where Configured	The full path of the file where the mount entry is stored.

6. Fill in the fields identifying the directories to be mounted. The information is similar to the fields for standard mount; see Step 5 in “Standard-Mount a Shared File System Using Text-Based HP SMH” (page 80). For details of the other options, see the *automount(1M)* manpage.

Standard-Mount a Shared File System Using HP-UX Commands

1. Log in to the *client* as superuser.
2. Ensure the client is configured to mount file systems via NFS. The simplest method is to use HP SMH; see “Enabling and Disabling the NFS Server and Client” (page 76).
3. Create the local directory on the client if it does not exist, for example:
`# mkdir /opt/adobe`



NOTE: If the directory does exist, its contents will be hidden when you mount the remote directory, and will not be usable until you unmount it.

4. Add an entry to /etc/fstab so the file system will be automatically mounted at boot time.

```
nfs_server:/nfs_server_dir /client_dir nfs options 0 0
```

For example:

```
fancy:/opt/adobe /opt/adobe nfs defaults 0 0
```

5. Mount the remote file system.

The following command forces the system to reread /etc/fstab and mount all the file systems:

```
# /usr/sbin/mount -a
```

Or you can just mount the one file system:

```
# /usr/sbin/mount /opt/adobe
```

Troubleshooting NFS

Table 5-2 outlines some troubleshooting techniques for common NFS problems.

Table 5-2 Troubleshooting NFS

Problem	What To Do
Individual client can't import from one or more servers	<p>Verify the following on the <i>client</i>:</p> <ul style="list-style-type: none">The local directory exists on the client. If it does not exist, create it using <code>mkdir</code>. For example: <code># mkdir /opt/adobe</code>LAN cable is intact and connected, and all connections are live.<code>/etc/hosts</code> exists and has "Requisite Entries" (page 85).<code>/etc/fstab</code> exists and has "Requisite Entries" (page 85), and the entries still point to valid directories on the server.<code>/etc/resolv.conf</code> exists and has "Requisite Entries" (page 85) (DNS only)<code>/etc/rc.config.d/nfsconf</code> has <code>NFS_CLIENT=1</code> <p>View the file directly, or use HP SMH to see that <code>NFS_CLIENT</code> is enabled (see "Enabling and Disabling the NFS Server and Client" (page 76)).</p> <p>Verify the following on the <i>servers</i>:</p> <ul style="list-style-type: none">The directories the client is trying to mount exist and are listed in <code>/etc/dfs/dfstab</code>.The client has permission to mount them. <p>See Step 3 under "Share a Directory Using HP-UX Commands" (page 77).</p>
All clients can't import from a given server	<p>Do the following on the <i>server</i>:</p> <ul style="list-style-type: none">Ensure that the server system is up and running, and that the LAN connection between the server and clients is live (can you ping the clients from the server and vice versa?).Ensure that <code>/etc/rc.config.d/nfsconf</code> has <code>NFS_CORE=1</code>, and <code>NFS_SERVER=1</code> or use HP SMH to see if NFS Server is enabled (see "Enabling and Disabling the NFS Server and Client" (page 76)).Ensure that the file systems that the clients are trying to mount are listed in <code>/etc/dfs/dfstab</code>. Check <code>/etc/dfs/dfstab</code> directly or use HP SMH (see "Sharing an HP-UX Directory" (page 77)).Restart the NFS server. See "Enabling and Disabling the NFS Server and Client" (page 76).If these remedies fail, and the configuration looks good (all the tests above), then the server may not have booted correctly; try rebooting the server.

Table 5-2 Troubleshooting NFS (continued)

Problem	What To Do
Stale NFS file handle This is common on NFS clients after a server has crashed, or been rebooted before clients have unmounted NFS file systems, or after /etc/dfs/dfstab has been changed on the server.	On the clients: <ul style="list-style-type: none">Use the <code>rmsf</code> command with the <code>-x</code> and <code>-H</code> options to remove stale device special files. For details, see <code>rmsf(1M)</code>. or ... <ul style="list-style-type: none">Ensure that there are no open files in the affected file systems; then try unmounting and remounting them. Try this first if <code>/etc/dfs/dfstab</code> has been changed on the server (directly or via HP SMH). On the server: <ul style="list-style-type: none">Run: <code># /usr/sbin/shareall</code> Try this first if server has just rebooted.
On an NFS server, <code>umount</code> fails.	<ul style="list-style-type: none">Check that all files are closed in the file system to be unmounted, and that it is not anyone's working directory on the system (host) from which it is to be unmounted. Note that although <code>fuser(1M)</code> can be used to check for open files, it is not able to detect files in a different directory opened within an editor.Try this if the directory is shared: <code># /usr/sbin/unshare dir</code>

Requisite Entries

The following entries are required in `/etc/hosts`, `/etc/fstab`, and `/etc/resolv.conf`:

- `/etc/hosts`:
 - System host name and IP address, for example:
`12.0.14.123 fredsyst fredsyst.mysite.myco.com`
 - An entry similar to the following:
`127.0.0.1 localhost loopback # [no SMTP]`
- `/etc/fstab`:
 - For standard mounts, an entry for each imported file system. See “Standard-Mount a Shared File System Using HP-UX Commands” (page 82).
- `/etc/resolv.conf` (needed for Domain Name Service (DNS) only):
 - The name of the domain in which this system resides, for example:
`domain mysite.myco.com`
 - At least one name server, for example:
`nameserver 12.0.14.165`

Recovering Network Services after a Power Failure

This section describes how to troubleshoot problems you and your system users are likely to encounter when rebooting after a general power failure or outage. The examples assume you are using DNS (Domain Name Service).

Symptoms and Keywords

```
RPC_PROG_NOT_REGISTERED  
name_server  
rcmd: hostname: Unknown host  
rcmd:hostname: Not in database  
rcmd:hostname: Access denied
```

What To Do

A. When the Domain Name Server Goes Down

If a system powers up *before* the Domain Name Server does, the system will not find the name server and, when users tries to reach another system, they will get the message:

```
rcmd: hostname: Unknown host
```

The simplest solution is to reboot the system *after* the name server has been rebooted.

B. When a Client Can't Import Directories from a Server

Do the troubleshooting checks described under “[Troubleshooting NFS](#)” (page 84). If these fail, and the client is getting messages such as:

```
rcmd: hostname: Not in database  
rcmd: hostname: Access denied
```

then do the following procedure:

1. Log in to the *server* as superuser.
2. Start HP SMH, as described in “[Starting Text-Based HP SMH](#)” (page 30).
3. Enable or restart the NFS server, as described in “[Enabling and Disabling the NFS Server and Client](#)” (page 76).
4. Exit HP SMH.
5. Log in to the *client* as superuser.
6. Start HP SMH, as described in “[Starting Text-Based HP SMH](#)” (page 30).
7. Enable the NFS client, as described in “[Enabling and Disabling the NFS Server and Client](#)” (page 76).
8. Exit HP SMH.

Moving or Reusing a Shared Directory

If you rename an NFS-mounted directory, NFS clients must unmount and remount the imported directory before they can see the new contents.

For example, if a server is sharing /opt/myapp, and you move /opt/myapp to /opt/myapp.old then rebuild and repopulate /opt/myapp, all the NFS clients *must* unmount and remount the directory, for example (as superuser on each client):

```
# umount /opt/myapp  
# mount -a
```

Any client on which this is not done will continue to see the former contents of /opt/myapp, that is /opt/myapp.old.

You can encounter the same problem in a slightly different way when you reuse an LVM volume.

For example, suppose you unmount an obsolete file system named /projects from a file server named fp_server, and subsequently reuse the logical volume, mounting a file system /newprojects on it.

Any client that fails to unmount /projects will see the contents of fp_server:/newprojects, labeled /projects.

Configuring File Transfer Protocol (FTP)

File Transfer Protocol (FTP) is a mechanism for copying files from one system to another. These sections provide configuration procedures and troubleshooting information.

Configuring Anonymous FTP

Anonymous FTP allows users who do not have an account on a given system to send files to, and retrieve them from, that system.

Configuring Anonymous FTP with Text-Based HP SMH

1. Log in to the *server* as superuser.
2. Start HP SMH; see “Starting Text-Based HP SMH” (page 30).
3. Navigate to **Network Services**.
 - a. Press **n, Networking and Communications**.
 - b. Press **s, Network Services Configuration**.
 - c. Press **k, Network Services**.
4. Highlight **Anonymous FTP**.
5. Press **Tab A** to pull down the **Actions** menu.

6. Select **Enable**.

HP SMH executes commands that create a password entry in the appropriate files, creates the `ftp` user account, and sets up the needed files and programs.

The entry in `/etc/passwd` is usually:

```
ftp:*:500:1:Anonymous FTP user:/home/ftp:/usr/bin/false
```

The created files and directories are:

```
dr-xr-xr-x  6 root      other          96 Oct 29 21:48 /home/ftp
dr-xr-xr-x  2 root      other          96 Oct 29 21:48 /home/ftp/etc
-r-----r--  1 root      other        1272 Oct 29 21:48 /home/ftp/etc/passwd
-r-----r--  2 root      other        226 Oct 29 21:48 /home/ftp/etc/group
-r-----r--  2 root      other        226 Oct 29 21:48 /home/ftp/etc/logginggroup
dr-xr-xr-x  4 root      other          96 Oct 29 21:48 /home/ftp/usr
dr-xr-xr-x  2 root      other          96 Oct 29 21:48 /home/ftp/usr/bin
----x---x  1 root      other        479232 Oct 10 21:39 /home/ftp/usr/bin/ls
dr-xr-xr-x  2 root      other          96 Oct 29 21:48 /home/ftp/usr/lib
-r-----r--  1 root      other        17782 Oct 10 21:43 /home/ftp/usr/lib/tztab
dr-xr-xr-x  2 root      other          96 Oct 29 21:48 /home/ftp/dist
drwxrwxrwx  2 ftp       other          96 Oct 29 21:48 /home/ftp/pub
```

Configuring Anonymous FTP with HP-UX Commands

1. Add user `ftp` to `/etc/passwd`, usually:

```
ftp:*:500:1:anonymous FTP:/home/ftp:/usr/bin/false
```

The password field should be `*`, the group membership should be `guest`, or, as in this example, `other`, and the login shell should be `/usr/bin/false`.

In this example, user `ftp`'s user ID is 500, and the anonymous FTP directory is `/home/ftp`.

2. Create the `ftp` home directory:

- Create the `ftp` home directory that you referred to in the `/etc/passwd` file, usually:

```
# mkdir /home/ftp
```

- Create the subdirectories `usr/bin` and `usr/lib` under the `ftp` home directory, usually:

```
# cd /home/ftp
# mkdir -p usr/bin
```

3. Copy the `ls` command from `/sbin` to `/home/ftp/usr/bin`, and set the permissions on the command to execute only (mode 0111):

```
# cp /sbin/ls /home/ftp/usr/bin
# chmod u=x,g=x,o=x /home/ftp/usr/bin/ls
```

4. Set the owner of the `/home/ftp/usr/bin` and `/home/ftp/usr` directories to `root`, and set the permissions to read-execute (not writable) (mode 0555):

```
# chown root /home/ftp/usr/bin
# chmod u=rx,g=rx,o=rx /home/ftp/usr/bin
```

- ```
chown root /home/ftp/usr
chmod u=rx,g=rx,o=rx /home/ftp/usr
```
5. Create the subdirectory etc under the ftp directory:  

```
cd /home/ftp
mkdir etc
```
  6. Copy /etc/passwd and /etc/group to /home/ftp/etc.  
These files are required by the ls command, to display the owners of files and directories under /home/ftp.  

```
cp /etc/passwd /home/ftp/etc
cp /etc/group /home/ftp/etc
```
  7. In all entries in /home/ftp/etc/passwd, replace the password field with an asterisk (\*), and delete the shell field, for example:  

```
ftp:*:500:1:anonymous ftp:/home/ftp:
tom:*:8996:20::/home/tom:
```
  8. In all entries in /home/ftp/etc/group, replace the password field with an asterisk (\*):  

```
users:*:20:acb
guest:*:21:ftp1
```
  9. Change the owner of the files in /home/ftp/etc to root, and set the permissions to read only (mode 0444):  

```
chown root /home/ftp/etc
chmod u=r,g=r,o=r /home/ftp/etc
```
  10. Create a directory pub (for public) under /home/ftp, and change its owner to user ftp and its permissions to writable by all (mode 0777).  
Anonymous FTP users can put files in this directory to make them available to other anonymous FTP users.  

```
mkdir /home/ftp/pub
chown ftp /home/ftp/pub
chmod u=rwx,g=rwx,o=rwx /home/ftp/pub
```

You can create other directories to provide separate categories, such as /home/ftp/draft and /home/ftp/final.
  11. Create a directory dist (for distribution) under /home/ftp. Change its owner to root and its permissions to writable only by root (mode 0755).  
Anonymous FTP users can read but not alter these directories.  

```
mkdir /home/ftp/dist
chown root /home/ftp/dist
chmod u=rwx,g=rx,o=rx /home/ftp/dist
```

12. Change the owner of user `ftp`'s home directory to `root` and the permissions to not writable (mode 0555):

```
chown root /home/ftp
chmod u=rw,g=rw,o=rw /home/ftp
```

## Troubleshooting FTP Login

**Symptom:** Some or all users can't `ftp` to an HP-UX system.

If *no* users can `ftp` to a given system, check first of all that `inetd` is running on that system:

```
ps -ef | grep inetd
```

If `inetd` is not running, start it:

```
/usr/sbin/inetd
```

It is also possible that the FTP service is disabled. Check `/etc/inetd.conf` for the following line:

```
FTP stream tcp nowait root /usr/lbin/FTPD FTPD -l
```

If this line does not exist, or is commented out (preceded by a pound sign (#)) add it (or remove the pound sign) and restart `inetd`:

```
/usr/sbin/inetd -c
```

You can also use HP SMH to check for the status of FTP and enable it if necessary.

- On text-based HP SMH, navigate to **Networking and Communications**→**Network Services Configuration**→**Network Services**.
- On web-based HP SMH, navigate to **Tools**→**Network Services**→**Configuration****Network Services**.

## Setting Up `/etc/shells`

**Problem:** FTP calls `getusershell` which by default checks password information (that is, the entry in `/etc/passwd` for the user who is trying to log in) against a fixed list. If the shell isn't on the list, FTP won't let the user in, so if you use an unusual shell you may not be able to `ftp` even to your own system.

`getusershell` can be made aware of other shells via `/etc/shells`. Perform the following steps on the system that is rejecting FTP logins:

1. If necessary, update all the old-style shell entries in `/etc/passwd`.  
Convert all `/bin/shellname` to `/usr/bin/shellname`.
2. Create `/etc/shells` and list all the shells that appear in `/etc/passwd`.

For more information, see `getusershell(3C)` and `shells(4)`.

## Configuring HP-UX Systems for File Transfer

Transferring files between computers is a common workgroup activity. When you're mixing HP-UX systems and PCs in a workgroup, network transfers are usually the most efficient, and sometimes the *only*, way to transfer files from one type of system to another. Many HP-UX systems are not equipped with floppy disk drives, and many PCs are not equipped with DDS drives or other external file storage peripherals often found on HP-UX systems.

### FTP (File Transfer Protocol)

One of the utilities/protocols common to both Windows NT and HP-UX systems is FTP (file transfer protocol). FTP is a client/server protocol. The **FTP client** is the program you run on your local system to communicate with the **FTP server** on the remote system.

#### FTP Client Software

On HP-UX systems, the FTP client is the program `/usr/bin/ftp`. On Microsoft's Windows operating systems you start the FTP client by issuing the `ftp` command from the command prompt.

#### FTP Server Software

Shipped as part of Windows operating systems for PCs (but not necessarily installed initially) are a group of utilities collectively known as the "Microsoft Peer Web Services." One of the services in this collection is an "FTP publishing service" that enables you to `ftp` files to and from your PC while sitting at one of your HP-UX systems. This service is the FTP server that runs on your PC. On HP-UX systems, the FTP server is the `ftpd` daemon, started as needed by the `inetd` daemon when FTP requests come in from clients on other systems.

As the name implies, file transfer protocol is used to transfer files from one system to another. Transferring files from one computer to another is a two-stage process. You must first establish a connection with, and log in to, the remote computer; then, you must locate and transfer the files you want to move to or from the remote computer.

### Establishing an FTP Connection from HP-UX to a PC



**NOTE:** See also "Establishing an FTP Connection from a PC to HP-UX" (page 99).

Before starting the following procedure, make sure FTP is set up for the kind of access you need. The default is to allow only anonymous access. If you want to allow individual user access, you can do this with the Internet Services Manager on your PC.

1. On your HP-UX system, start the FTP utility by entering the command:  
`# /usr/bin/ftp`
2. Open a connection to your PC using `ftp`'s `open` command:

```
ftp> open vectrapc1.net2.corporate
```

If the connection is successful, FTP will let you know that you are connected and display information about the PC's FTP server:

```
Connected to vectrapc1.net2.corporate.
```

```
220 vectrapc1 Microsoft FTP Service (Version 2.0).
```

If your connection succeeded, proceed to Step 3. If it fails, use Table 5-3 (page 93).

**Table 5-3 Troubleshooting the FTP Connection to a PC**

| TROUBLESHOOTING INFORMATION                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>If the connection is <i>not</i> successful, FTP will let you know that the connection failed. The displayed error message will vary depending on what is the cause of the failed connection:</p> <p><input type="checkbox"/> ftp: connect: Connection refused<br/>The most likely cause of this message is:<br/>– <i>Problem:</i> The FTP publishing service on the Windows NT-based PC is not running (has not been started).<br/><i>Solution:</i> Start the FTP server on the PC.</p> <p><input type="checkbox"/> ftp: connect: Connection timed out<br/>Possible causes of this error message include:<br/>– <i>Problem:</i> Your PC is not currently running.<br/><i>Solution:</i> Make sure your PC is turned on, and running (the Windows NT operating system has been booted).<br/>– <i>Problem:</i> Your PC is not currently reachable on the network.<br/><i>Solution:</i> Make sure that the your PC is physically connected to the network and that there are no network outages or breaks between your PC and your HP-UX system.</p> <p><input type="checkbox"/> ftp: vectrapcl: Unknown host<br/>Possible causes of this error message include:<br/>– <i>Problem:</i> You typed the name of your PC incorrectly.<br/><i>Solution:</i> Verify that you entered the name of your PC correctly in the open command. Depending on where in your network structure the PC is located with respect to your HP-UX system, it might be necessary to fully qualify the PC name. For example:<br/><code>ftp&gt; open vectrapcl</code><br/>is probably sufficient if your PC is on your local network segment, but a more fully qualified name, for example:<br/><code>ftp&gt; open vectrapcl.net2</code><br/>or<br/><code>ftp&gt; open vectrapcl.net2.corporate</code><br/>will likely be needed to access your PC if it is located elsewhere in your network (across a router or gateway). If all of the above fail, try using the IP address of the PC in place of the name. For example:<br/><code>ftp&gt; open 15.nn.xx.2</code><br/>– <i>Problem:</i> Your PC is not formally known to your network<br/><i>Solution:</i> Make sure that networking services, particularly TCP/IP services have been properly configured on your Windows NT operating system. The computer must have its own valid IP address, and you must assign it a DNS host name and domain. These are assigned via the Network service in the Windows NT Control Panel.</p> |

### 3. Enter login information

When you have successfully connected to your PC, another message will follow the Connected to... message:

Name (vectrapc1.net2.corporate:userx) :

This message is actually a login prompt, and there are several ways to respond to it:

- **Press Return to accept the default response.**

In the above example, there are three parts to the displayed prompt:

1. The word Name
2. The network name for your PC (vectrapc1.net2.corporate)
3. The default user name (userx); this is usually the name of the HP-UX account that you were using when you issued the `ftp` command in Step 1.

If you press **Return**, `ftp` will attempt to log you in to the PC using the same name as you used to log into HP-UX. You will then be prompted to enter your password. If, after noting the following caution and you feel comfortable doing so, enter the password.



**CAUTION:** It is important to note here that any characters you type at your keyboard, including your user name and password will be transmitted over the network to your PC *unencrypted*.

Although it is unlikely, especially if your network is strictly an internal network, it is possible that someone could be eavesdropping on your network lines and obtain your login information. If this is a concern to you, HP strongly recommends that you use the anonymous login option described in the following text.

---

- **Enter a valid account name and password for your PC.**

If the PC account you want to log in to is different from the user name you used to log in to HP-UX, enter the user name for the PC account at the prompt. You will then be prompted to enter the password for the account. If, after noting the preceding caution and you feel comfortable doing so, enter the account's password.

- **Use FTP's anonymous login feature.**

Because account names and passwords that you enter from the keyboard during the FTP login process are sent to the remote computer unencrypted (making this sensitive information vulnerable to network eavesdroppers), FTP provides a way to access a remote computer using what is known as an anonymous login. To use this feature, enter the word `anonymous` at the prompt:

```
Name (vectrapc1.net2.corporate:userx): anonymous
```

You will then be prompted to enter a password in a special way:

```
331 Anonymous access allowed, send identity (e-mail name) as password.
```

Instead of entering the actual password for an account, enter your e-mail address as a way of identifying yourself to the FTP server:

```
Password: userx@net2.corporate
```

After successfully entering the PC account information you will be logged in to the PC and placed in the directory designated as the *ftp-root* directory in your Windows NT configuration.

Using the FTP client's cd command, remote users of the PC can access:

- The *ftp-root* directory
- Any of the subdirectories of the *ftp-root* directory
- Selected other directories on the PC that have specifically been made available by the administrator of the PC

For information about how to make those other directories available, refer to the online documentation associated with the Microsoft Internet Service Manager.

### On the HP-UX System: Retrieving a File from the PC

Once you have made a connection and logged in to the PC from your HP-UX system (See "Establishing an FTP Connection from HP-UX to a PC" (page 91)), you are ready to retrieve a file from the PC.

1. Locate the file you want to retrieve from your PC. You can use FTP's cd and ls commands pretty much as you would in an HP-UX shell (sh, ksh, csh, etc.). If it is not in the PC's *ftp-root* directory, use FTP's change directory command (cd) to move to the directory on the PC where the file exists.

2. Determine whether the file you are trying to transfer is an ASCII (text) file or a binary (non-ASCII) file and set the transfer mode accordingly:

- a. For ASCII files, set the transfer mode using FTP's `ascii` command:

```
ftp> ascii
```

This enables character conversions such as end-of-line carriage return stripping to occur.

- b. For binary files (graphics files, sound files, data base files, etc.), set the transfer mode using FTP's `binary` command:

```
ftp> binary
```

This causes FTP to use an eight-bit-wide (byte) transfer rather than a seven-bit-wide (character) transfer. This is very important as most non-ASCII formats are dependent on that eighth bit of each byte. *Your binary files will be corrupted if you transfer them using ASCII mode.*



**TIP:** If you are unsure of the format of the file you are transferring (ASCII or binary) set the file type to binary. ASCII files will not be corrupted if transferred in binary mode; however, end-of-line character stripping will not occur.

3. Transfer the file using FTP's `get` command.

### Example 5-1 Retrieve an ASCII File with FTP

To retrieve the ASCII file `phone.dat` (located in the subdirectory called `data`, under the `ftp-root` directory) from the PC:

```
ftp>cd data
ftp>ascii
ftp>get phone.dat
```

### Example 5-2 Retrieve a Binary File with FTP

To then retrieve the graphics file `net2.jpg` from the subdirectory called `pics` (located under the `ftp-root` directory):

```
ftp>cd ../pics
ftp>binary
ftp>get net2.jpg
```

## On the HP-UX System: Sending a File to the PC

Once you have made a connection and logged in to the PC from your HP-UX system (See "Establishing an FTP Connection from HP-UX to a PC" (page 91)), you are ready to transfer a file to the PC.

1. Locate the file you want to send. You can use FTP's `lcd` and `!` (execute a local shell command) commands to locate the file on your local system if it is not in the directory that was your current working directory at the time you started `ftp`. Also, if the file is not in your current directory, you can specify a full (absolute) path name for the file you want to send to your PC.

2. Determine whether the file you are trying to transfer to your PC is an ASCII text file or a binary (non-ASCII) file and set the transfer mode accordingly:

- a. For ASCII (plain text) files, set the transfer mode using FTP's **ascii** command:

```
ftp>ascii
```

This enables character conversions such as those that handle the differences between how the ends of lines are handled between differing types of operating systems.

- b. For binary files (graphics files, sound files, data base files, etc.), set the transfer mode using FTP's **binary** command:

```
ftp>binary
```

This causes FTP to use an eight-bit-wide byte transfer rather than a seven-bit-wide character transfer. This is very important as most non-ASCII formats are dependent on that eighth bit of each byte. *Your binary files will be corrupted if you transfer them using ASCII mode.*



**TIP:** If you are unsure of the format of the file you are transferring (ASCII or binary), set the file type to **binary**. ASCII files will not be corrupted if transferred in binary mode; however, end-of-line character handling will not occur.

---

3. Transfer the file using FTP's **send** command.

### **Example 5-3 Send from Different Directory**

---

To send the ASCII file `phone.dat` (located in the `/var/tmp` directory on your HP-UX system) to the PC:

```
ftp>lcd /var/tmp
ftp>ascii
ftp>send phone.dat
```

— OR —

```
ftp>ascii
ftp>send /var/tmp/phone.dat
```

---

### **Example 5-4 Send from Current Directory**

---

To send the graphics file `roadmap.jpg` from the current working directory:

```
ftp>binary
ftp>send roadmap.jpg
```

---



**NOTE:** See also “Establishing an FTP Connection from HP-UX to a PC” (page 91). If you have a third-party program, use those instructions instead.

1. On your PC, start the FTP utility:
  - a. Click the Start bar in the lower-left corner of your PC’s screen.
  - b. Click Programs in the pop-up menu.
  - c. Click Accessories in the next pop-up menu.
  - d. Click Command Prompt in the final pop-up menu.
  - e. Type `ftp` at the prompt in the window.
2. Open a connection to your HP-UX system using FTP’s open command:

```
ftp> open flserver.net2.corporate
```

If the connection is successful, FTP will let you know that you are connected and display information about the FTP server on the HP-UX system:

```
Connected to flserver.net2.corporate.
220 flserver FTP Server (Version 1.7.111.1) ready.
```

If your connection succeeded, proceed to Step 3.

If the connection is *not* successful, FTP will let you know that the connection failed. The displayed error message will vary depending on what is the cause of the failed connection:

- `ftp: connect: Connection refused`

Possible causes of this error message include:

- *Problem:* The internet daemon (`inetd`) is not running on your HP-UX system.

*Solution:* The real problem is that the `ftpd` daemon is not running, but it is usually `inetd` that starts `ftpd` on an as-needed basis. `inetd` is usually started up when you boot your computer. If your HP-UX system is in single-user mode you will need to switch it to a run-level of 2 or higher.

- *Problem:* The FTP daemon (`ftpd`) is not running.

*Solution:* Verify that there is a valid entry in the file `/etc/inetd.conf` for the `ftpd` daemon. The entry should look like this:

```
ftp stream tcp nowait root /usr/lbin/ftpd ftp -lconf
```

Make sure that the entry is not commented out (no # in the first column).

Make the appropriate repairs and use the command

```
/usr/sbin/inetd -c
```

to have `inetd` reread its configuration file.

- `ftp: connect: Connection timed out`

Possible causes of this error message include:

- *Problem:* Your HP-UX system is not currently running.

*Solution:* Make sure your HP-UX system is turned on, and running (the system has been booted).

- *Problem:* Your HP-UX system is not currently reachable on the network.

*Solution:* Make sure that the your HP-UX system is physically connected to the network and that there are no network outages or breaks between your PC and your HP-UX system.

- `ftp: flserver: Unknown host`

Possible causes of this error message include:

- *Problem:* You typed the name of your HP-UX system incorrectly.

*Solution:* Verify that you entered the name of your HP-UX system correctly in the `open` command. Depending on where in your network structure the system is located with respect to your PC, it might be necessary to fully qualify the HP-UX system name. For example:

`ftp>open flserver`

is probably sufficient if your PC is on your local network segment, but a more fully qualified name, for example:

`ftp>open flserver.net2`

or

`ftp>open flserver.net2.corporate`

will likely be needed to access your HP-UX system if it is located elsewhere in your network (across a router or gateway). If all of the above fail, try using the IP address of the HP-UX system in place of the name. For example:

`ftp>open 15.nn.xx.100`

- *Problem:* Your HP-UX system is not formally known to your network.

*Solution:* Make sure that networking services, particularly TCP/IP services have been properly configured on your HP-UX system. The computer must have its own, valid IP address, and you must assign it a valid host name.

### 3. Enter login information

When you have successfully connected to your HP-UX system, another message will follow the Connected to... message:

```
Name (flserver.net2.corporate:(none)):
```

This message is actually a login prompt, and there are two ways to respond to it:

- **Enter a valid account name and password for your PC.**

You will then be prompted to enter the password for the account. If after noting the following caution you feel comfortable doing so, enter the account's password.



**CAUTION:** It is important to note here that any characters you type at your keyboard, including your user name and password will be transmitted over the network to your PC *unencrypted*!

Although it is unlikely, especially if your network is strictly an internal network, it is possible that someone could be eavesdropping on your network lines and obtain your login information. If this is a concern to you, HP strongly recommends that you use the anonymous login option described in the following text.

- **Use FTP's anonymous login feature**

Because account names and passwords that you enter from the keyboard during the FTP login process are sent to the remote computer unencrypted (making this sensitive information vulnerable to network eavesdroppers), FTP provides a way to access a remote computer using what is known as an anonymous login. To use this feature, enter the word anonymous at the prompt:

```
Name (flserver.net2.corporate:userx):anonymous
```

You will then be prompted to enter a password in a special way:

```
331 Anonymous access allowed, send identity (e-mail name) as password.
```

Instead of entering the actual password for an account, enter your e-mail address as a way of identifying yourself to the FTP server:

```
Password: glenda@net2.corporate
```

After successfully entering the HP-UX account information you will be logged in to your HP-UX system and placed in the directory designated as the *ftp-root* directory.

Using the FTP client's cd command, remote users (logged in anonymously) can access:

- the *ftp-root* directory
- any of the subdirectories of the *ftp-root* directory

## On the PC: Retrieving a File from the HP-UX System

Once you have made a connection and logged in to your HP-UX system from your PC (See “Establishing an FTP Connection from a PC to HP-UX” (page 99)) you are ready to retrieve a file from the HP-UX system.

1. Locate the file you want to retrieve from your HP-UX system. You can use FTP’s `cd` and `ls` commands pretty much as you would in an HP-UX shell (`sh`, `ksh`, `csh`, etc.). If it is not in the home directory for the HP-UX account that you logged in to, use FTP’s change directory command (`cd`) to move to the directory on the HP-UX system where the file exists.

2. Determine whether the file you are trying to transfer is an ASCII file or a binary (non-ASCII) file and set the transfer mode accordingly:
  - a. For ASCII (plain text) files, set the transfer mode using FTP's `ascii` command:  
`ftp>ascii`  
This enables character conversions such as end-of-line carriage return stripping to occur.
  - b. For binary files (graphics files, sound files, database files, etc.), set the transfer mode using FTP's `binary` command:  
`ftp>binary`

This causes FTP to use an eight-bit-wide (byte) transfer rather than a seven bit wide (character) transfer. This is very important as most non-ASCII formats are dependent on that eighth bit of each byte!

---



**CAUTION:** Your binary files will be corrupted if you transfer them using ASCII mode.

---



**TIP:** If you are unsure of the format of the file you are transferring (ASCII or binary) set the file type to `binary`. ASCII files will not be corrupted if transferred in binary mode, however end-of-line character stripping will not occur.

---

3. Transfer the file using FTP's `get` command.

#### **Example 5-5 Retrieve an ASCII File with FTP**

---

To retrieve the ASCII file `phone.dat` (located in the subdirectory called `data`, under the `home` directory for your account) from the HP-UX system:

```
ftp>cd data
ftp>ascii
ftp>get phone.dat
```

---

#### **Example 5-6 Retrieve a Binary File with FTP**

---

To then retrieve the graphics file `net2.jpg` (from the subdirectory called `pics` located under the `home` directory):

```
ftp>cd ../pics
ftp>binary
ftp>get net2.jpg
```

---

## On the PC: Sending a File to the HP-UX System

Once you have made a connection and logged in to your HP-UX system (See “Establishing an FTP Connection from a PC to HP-UX” (page 99)), you are ready to transfer a file to the your HP-UX system.

1. On your PC, locate the file you want to send. You can use FTP’s `lcd` and `!commands` to locate the file on your local system if it is not in the directory that was your current working directory at the time you started `ftp`. If the file is not in your current directory, you can specify a full (absolute) path name for the file you want to send to your HP-UX system, or use FTP’s `lcd` command to move to the directory containing the file.

2. Determine whether the file you are trying to transfer to your HP-UX system is an ASCII file or a binary (non-ASCII) file and set the transfer mode accordingly:
  - a. For ASCII (plain text) files, set the transfer mode using FTP's `ascii` command:

```
ftp>ascii
```

This enables character conversions such as those that handle the differences between how the ends of lines are handled between differing types of operating systems.

- b. For binary files (graphics files, sound files, database files, etc.), set the transfer mode using FTP's `binary` command:

```
ftp>binary
```

This causes FTP to use an eight-bit wide (byte) transfer rather than a seven bit wide (character) transfer. This is very important as most non-ASCII formats are dependent on that eighth bit of each byte! *Your binary files will be corrupted if you transfer them using ASCII mode.*



**TIP:** If you are unsure of the format of the file you are transferring (ASCII or binary) set the file type to binary. ASCII files will not be corrupted if transferred in binary mode, however end-of-line character handling will not occur.

3. Transfer the file using FTP's `send` command.

#### **Example 5-7 Send an ASCII File with FTP**

To send the ASCII file `phone.dat` (located in the `C:\office_stuff` directory on your PC) to your HP-UX system:

```
ftp>lcd C:\office_stuff
ftp>ascii
ftp>send phone.dat
— OR —
ftp>ascii
ftp>send C:\office_stuff\phone.dat
```

#### **Example 5-8 Send a Binary File with FTP**

To send the graphics file `roadmap.jpg` from the current working directory:

```
ftp>binary
ftp>send roadmap.jpg
```

# Interfacing with Microsoft Windows Systems

The following topics describe the process of adding PC/NT systems into a work group.

- “Hardware Connections” (page 106)
- “Configuring HP-UX Systems for Terminal Emulation” (page 107)
  - “Telnet” (page 107)
  - “Other Terminal Emulators” (page 110)
- “Configuring HP-UX Systems for File Transfer” (page 91)
  - “FTP (File Transfer Protocol)” (page 91)
- “Sharing Directories between HP-UX and MS Windows” (page 106)

## Hardware Connections

Adding a personal computer (PC) to a workgroup is much more a logical operation than a physical one. The only requirement from a hardware perspective is to give the personal computer physical access to the other computers in the workgroup. This connection is usually (but not always) a network connection. It could, however, be a modem (dial-in) connection: a telephone-based UUCP connection, or a Serial Line Internet Protocol (SLIP) connection for example.

The requirements of this connection depend on how you plan to interact with the PC. For example, occasionally transferring small ASCII files or exchanging text-based e-mail between the users of the PC and the users of your HP-UX computers isn't likely to be a problem for a serial line because comparatively little data are being transferred between computers. However, if you plan to constantly share X Window screens between the HP-UX systems and the PC, you had better have a high-speed connection, such as a network connection between the two types of computers, or the performance of your applications will be unacceptably slow (if they work at all).

When connecting the PC to your other computers, you should consider:

- The amount of data to be exchanged between the PC and the other computers in your workgroup
- How often you plan to access the data on the PC (occasionally? frequently? constantly?)
- The type of data you want to exchange (ASCII text? graphics? sound? video?)
- How will you exchange the data (file transfer? shared windowing environment? electronic mail?)

## Sharing Directories between HP-UX and MS Windows

You can use the HP CIFS product (or third-party products) to share data between HP-UX systems and Windows systems.

## HP CIFS

HP CIFS provides HP-UX with a distributed file system based upon Microsoft's CIFS (Common Internet File System) protocol, also known as the SMB (Server Message Block) protocol. The SMB protocol is the native file-sharing protocol in Microsoft Windows and OS/2 operating systems and is the standard way that millions of PC users share files across corporate intranets.

HP CIFS implements both the server and client components of the CIFS protocol on HP-UX. This means that HP-UX file systems can be mounted onto Windows systems and Windows file systems can be mounted onto HP-UX systems.

The HP CIFS Server is based upon Samba and provides file as well as print services to CIFS clients including Windows NT, XP, 2000, Server 2008, and Vista; and other HP-UX machines running the HP CIFS Client software.

The HP CIFS Client enables HP-UX users to mount as UNIX file systems PC shares from CIFS file servers including Windows servers and HP-UX machines running the HP CIFS Server software. The HP CIFS client also offers an optional Pluggable Authentication Module (PAM) that implements the Windows NTLM authentication protocols. When installed and configured within HP-UX's PAM facility, this allows HP-UX users to be authenticated against a Windows authentication server.

For detailed information on how to install, configure and use the HP CIFS server and client software, see the *HP CIFS Server Administrator's Guide* and the *HP CIFS Client Administrator's Guide*, both available at <http://docs.hp.com>.

## Configuring HP-UX Systems for Terminal Emulation

The primary reason for having a computer in a workgroup (regardless of what type of computer it is) is so that its users can access the resources of other computers in the workgroup.

A common way to access the resources of another computer is to log into the remote computer using a terminal emulation program such as Telnet.

### Telnet

The `telnet` utility is a standard part of the HP-UX operating system, and a Telnet *client* is included in versions of Microsoft's Windows operating systems. It is used to log in to a remote system from a personal computer (PC) or an HP-UX system.

The remote system can be a UNIX-based system (such as an HP-UX system), or a PC running Telnet server software. Initially, Windows includes a Telnet *client* program, which can be used to log in to remote computers, but does not include a Telnet *server* application, which would allow other computers to "Telnet in" to a Windows operating system. On HP-UX systems, the Telnet server software is known as the `telnetd` daemon.

## Using Telnet to Log in to a PC from an HP-UX System

To use Telnet to log in to a personal computer from your HP-UX system, you will need to:

1. Make sure that the PC is running, and reachable via your network.
  - a. Turn on the PC and boot up the Windows NT operating system.
  - b. Make sure that your PC has networking services configured, and has a network address (IP Address).
2. Make sure that the PC is running Telnet server software.
  - a. Install a version of Telnet server software.



**NOTE:** Microsoft's Windows operating systems do not initially include Telnet *server* software. Commercial and shareware versions of Telnet server software are available from a variety of sources.

- b. Configure, and start the Telnet server software according to the instructions that come with it.
3. On your HP-UX system, start the `telnet` utility and open a connection to the PC you are trying to access. For example:

```
$ /usr/bin/telnet
telnet>open vectrapc1.net2.corporate
Trying...
Connected to vectrapc1.net2.corporate.
Escape character is '^]'.
Local flow control off
```

(A pleasant telnet server/OS identification message)

login:



**TIP:** You can shorten the connection process by using `telnet` in noninteractive mode. To do this, specify the name of the PC that you are trying to connect to as an argument on the command line when you start up `telnet`. For example:

```
$ /usr/bin/telnet vectrapc1.net2.corporate
```

4. Log in using the same user name and password as you would if you were sitting at the PC's keyboard. How you specify the NT domain information will vary depending on the Telnet server software that you are using. Follow the instructions that come with your Telnet server software or the prompts that the server software gives you during the login process.

## Using Telnet to Log in to an HP-UX System from a PC

1. Make sure that the PC is running, and reachable via your network.
  - a. Turn on the PC and boot up the Windows NT operating system.
  - b. Make sure that your PC has networking services configured, and has a network address (IP address).
2. Make sure that the `telnethd` daemon is running on your HP-UX system.

The `telnethd` daemon is not usually run directly. Copies of `telnethd` are started by the `inetd` daemon when requests arrive over the network for Telnet services. Therefore:

- a. Verify that an entry for `telnethd` exists in the configuration file `/etc/inetd.conf`; the entry should look like this:

```
telnet stream tcp nowait root /usr/lbin/telnethd telnetd
```

- b. Verify that the file `/etc/services` has an entry that looks like this:

```
telnet 23/tcp # Virtual Terminal Protocol
```

- c. Verify that the `inetd` daemon is running. On a networked system running at or above run level 2, `inetd` is automatically started by the script `/sbin/rc.2.d/S500inetd` during the boot-up sequence. You can verify that it is running by issuing the following command:

```
/usr/bin/ps -ef | grep inetd
```

3. On your PC, start the Telnet *client* software.

If you are using the Telnet client that comes with the Windows operating system, you can start the client by:

- a. Clicking on the Start bar in the lower-left corner of your PC's screen
- b. Clicking Programs in the resulting pop-up menu
- c. Clicking Accessories in the resulting pop-up menu
- d. Clicking on Telnet in the final pop-up menu

4. Use the Telnet client to connect to your HP-UX system.

If you are using the Telnet client software that comes with the Windows operating system, you can connect to your HP-UX system by:

- a. Clicking on the Connect menu item in the upper-left corner of your `telnet` window.
- b. Clicking on the Remote System . . . menu item from the connect menu.
- c. Entering the name of your HP-UX system in the Host Name field of the resulting dialog box (leave the Port field set to `telnet`).
- d. Clicking on the Connect button in the lower-left corner of the dialog box.

## Other Terminal Emulators

Telnet is only one of many terminal emulators — sometimes known as virtual terminals — that can be used to log in to remote systems, but in the UNIX world it is a common one.

Another that is often supported by software packages on the PC for interacting with UNIX systems is `rlogin`. The `rlogin` daemon on HP-UX systems is `rlogind`. The setup and use of `rlogin` between HP-UX systems and PCs is quite similar to that for Telnet, especially on the HP-UX end. `rlogin` (client or server) software is not part of Windows operating systems as originally shipped; however, commercial and shareware versions of `rlogin` can be found for your Windows NT-based PCs.



**IMPORTANT:** The `telnet` and `rlogin` terminal emulators do not provide secure communications. For more secure communications, you can use a secure shell (`ssh`) client to access HP-UX. For details on using a secure shell, see `ssh(1)`.

---

# 6 Configuring Printers

Printers can be connected in three ways:

- Local: physically connected to the system.
- Remote: physically connected to another system.
- Network: installed as a domain on a LAN.

This chapter describes how to perform the following procedures:

- “Starting and Stopping the LP Spooler” (page 111)
- “Adding a Local Printer to the LP Spooler” (page 112)
- “Adding a Remote Printer to the LP Spooler” (page 113)
- “Adding a Network Printer with HP JetDirect” (page 116)
- “Creating a Printer Class” (page 117)
- “Removing a Printer from the LP Spooler” (page 118)
- “Removing a Printer from a Printer Class” (page 119)
- “Removing a Printer Class” (page 120)



**NOTE:** Configured printers can be managed with the LDAP-UX Printer Configurator Services. See the *LDAP-UX Client Services B.04.00 Administrator’s Guide* for details.

For conceptual information about print-management topics, see the *HP-UX System Administrator’s Guide: Overview*. Such topics include:

- Planning your printer configuration.
- Remote spooling concepts.
- Printer model files: providing information for the -m option of the lpadmin command. See also *lpadmin(1M)*.
- Printer queues and printer classes for grouping printers logically.
- A processing overview.

For procedures on maintaining your printer environment, see the *HP-UX System Administrator’s Guide: Routine Tasks*. Such topics include:

- Controlling the flow of print requests to printer queues with accept and reject. See also *accept(1M)*.
- Starting and stopping locally queued print jobs from being sent to the associated printer with enable and disable. See also *enable(1)*.

## Starting and Stopping the LP Spooler

Before you can print using the LP spooler, one or more printers must be configured and the scheduler started. At the first system boot, the scheduler is not started if no printers are configured.

To add a printer to the spooler, see “Adding a Local Printer to the LP Spooler” (page 112), “Adding a Remote Printer to the LP Spooler” (page 113) or “Adding a Network Printer with HP JetDirect” (page 116).

## Starting and Stopping Spooler Using Text-Based HP SMH

1. As a privileged user, start text-based HP SMH. See “Starting Text-Based HP SMH” (page 30). You can invoke it with or without the X Window interface.
2. Press **p, Printers and Plotters**.
3. Highlight and select **Printers and Plotters**.  
An X Window or graphical text screen is displayed.
4. From the **Action** pulldown menu,
  - Choose **Stop Print Spooler** to stop the spooler.
  - Choose **Start Print Spooler** to start the spooler.

## Starting and Stopping Spooler Using HP-UX Commands

- To start the LP spooler:  
`# /usr/sbin/lpsched`
- To stop the LP spooler:  
`# /usr/sbin/lpshut`

## Adding a Local Printer to the LP Spooler



**NOTE:** Do not confuse adding a printer to the LP spooler with adding a printer to your system. Adding a printer to the LP spooler involves configuring the LP spooler. Adding a printer to your system involves connecting the printer to your computer and configuring the needed drivers in the kernel. For information on the latter, see the *Interface Card OL\* Support Guide*.

### Adding a Local Printer Using HP SMH

The easiest way to add a local printer to the LP spooler is to run HP SMH. HP SMH will also do some of the CDE configuration (if CDE is being used).

### Adding a Local Printer Using HP-UX Commands

1. Ensure that you have superuser capabilities.
2. Add the printer to the LP spooler. For example:  
`# /usr/sbin/lpadmin -plocal_printer -v/dev/lp -mHP_model -g7`  
See *lpadmin(1M)* for details on the options. Also see the *HP-UX System Administrator's Guide: Overview* for model information for the `-m` option.

3. If the printer being added will be the default printer, execute the following:

```
/usr/sbin/lpadmin -dlocal_printer
```

Allow print requests to be accepted for the newly added printer. For example:

```
/usr/sbin/accept local_printer
```

4. Enable the newly added printer to process print requests. For example:

```
/usr/bin/enable local_printer
```

5. Start the LP spooler if it is not already running:

```
/usr/sbin/lpsched
```

6. Test the printer using the LP spooler, then check the LP spooler's status. For example:

```
lp -dlocal_printer /etc/motd
```

```
lpstat -t
```

## Adding a Remote Printer to the LP Spooler

The easiest way to add a printer to a remote system is to run HP SMH. If you elect to use HP-UX commands, review Step 5 under "Adding a Remote Printer Using Text-Based HP SMH", as this information will also be required when performing the task manually.

### Adding a Remote Printer Using Text-Based HP SMH

---



**NOTE:** HP SMH does not verify that an actual printer exists on a remote system. Be sure the printer is installed and configured, and if necessary, use HP SMH to configure it on the remote system before adding it as a remote printer.

1. As a privileged user, start text-based HP SMH. See "Starting Text-Based HP SMH" (page 30). You can invoke it with or without the X Window interface.

2. Press **p, Printers and Plotters**.

3. Highlight and select **Printers and Plotters**.

An X Window or graphical text screen is displayed.

4. From the **Action** pulldown menu, choose **Add Remote Printer/Plotter**.

5. Provide information for the following data fields:

- **Printer Name**

The name you will use on this system.

- **Remote System Name**

The full domain name of the system that hosts the printer.

- **Remote Printer Name**

The name of the printer as it is known on the remote system.

- **Remote Printer is on a BSD system**  
Check the box if the remote operating system is BSD.
  - **Remote Cancel Name**  
The default is `rcmodel`.
  - **Remote Status Name**  
The default is `rsmodel`.
  - **Default Request Priority**  
This is the minimum priority (fence) that a print job must have in order to print on this printer. 0 is the lowest.
  - **Allow Anyone to Cancel a Request**  
Check the box if you want to allow anyone, not just the owner or a superuser, to cancel print requests for this printer. This can be a convenient idea for a small workgroup.
  - **Make this Printer the Default Destination**  
Check the box to make this printer your system's default.
6. When all fields are filled in, select **OK**. If the configuration was unsuccessful, HP SMH returns with troubleshooting information. Most likely problems will be related to the remote system configuration. Check as follows:
- a. Edit `/etc/services` (on the remote system), and if necessary, uncomment the line beginning with `printer` by removing the #.
  - b. Ensure no systems are restricted from access by `/var/adm/inetd.sec`.
  - c. Make sure `rlpdaemon` is running.

## Adding a Remote Printer Using HP-UX Commands

1. Ensure that you have superuser capabilities.
  2. Add the remote printer.
    - If the remote printer is on an HP-UX system, enter:

```
lpadmin -plocal_printer -v /dev/null -mrmodel \
> -ormremote_machine -orpremote_dest -ocmrcmodel \
> -osmrsmode1
```
    - If the remote printer is *not* on an HP-UX system, enter:

```
lpadmin -plocal_printer -v /dev/null -mrmodel \
> -ormremote_machine -orpremote_dest -ocmrcmodel \
> -osmrsmode1 -ob3
```
- See `lpadmin(1M)` for details on the options. Also see the *HP-UX System Administrator's Guide: Overview* for model information for the `-m` and `-o*` options.

3. Allow print requests to be accepted for the newly added remote printer. For example:  

```
/usr/sbin/accept local_printer
```
4. Enable the newly added printer to process print requests. For example:  

```
/usr/bin/enable local_printer
```
5. If the printer being added will be the default printer, execute the following:  

```
/usr/sbin/lpadmin -dlocal_printer
```
6. Start the LP spooler if it is not already running to process print requests.  

```
/usr/sbin/lpsched
```
7. Send a sample print job to the printer.
  - If it prints, the remote printing daemon (`rlpdaemon`) is active on the system and your task is completed.
  - If your print job does not print, the remote printing daemon (`rlpdaemon`) is not active yet on the remote machine. Activate the `rlpdaemon` on the host system where the remote printer resides, as follows:
    - Examine the file `/etc/inetd.conf` and look for the following line:  

```
printer stream tcp nowait root /usr/sbin/rlpdaemon rlpdaemon -i
```

If a # sign appears at the beginning of the line, the `rlpdaemon` line is commented out, preventing the printer from printing remotely.  
Edit the file `/etc/inetd.conf` to remove the # sign. Save the file.
    - Check `/etc/services` and look for:  

```
printer 515/tcp spooler #remote print spooling
```

If a # sign appears at the beginning of the line, the service is commented out, preventing the remote print spooler from serving the printer.  
Edit the file to remove the # sign in the first column. Save the file.
    - Reconfigure the Internet daemon `inetd`, forcing it to reread the `/etc/inetd.conf` file. Invoke the following command:  

```
/usr/sbin/inetd -c
```

Also, look for entries in `/var/adm/inetd.sec` that restrict which systems can send remote print requests.
8. Test the printer using the LP spooler, then check the LP spooler's status. For example:  

```
lp -dlocal_printer /etc/motd
lpstat -t
```

## Adding a Network Printer with HP JetDirect

The HP JetDirect software must be installed on your system. With HP JetDirect, printers can connect directly to the network. The printer uses a LAN connection and the HP JetDirect software transmits print requests.

For more information, see the *HP JetDirect Network Interface Configuration Guide*. Follow the instructions shipped with your printer or the network interface card for the printer. See also the *jetadmin(1)* manpage.

To add a network-based printer that uses the HP JetDirect Network Interface, you will need the following:

- The printer's full domain name or its Internet address.
- The local name that the LP spooler will use to refer to the printer.

## Adding A Network Printer Using Web-Based HP SMH

To start JetDirect from web-based HP SMH:

1. Ensure that an X Window System server is running on your local system. See “Using the X Window System” (page 32).
2. Start HP SMH from your local system. See “Starting an HP SMH Web Client” (page 24).
3. Navigate to **Tools**→**Printer Management**→**Configure Printers or Plotters**.
4. Click the **Run** button.

The **Printers and Plotters** X Window screen is displayed.

5. Select **Actions**→**Add Network-Based Printer/Plotter**→**Add Printer/Plotter Connected to HP JetDirect** and click **OK** to run the command /opt/hpnp/admin/jetadmin.
6. Enter 3 to choose option **3) Add printer to local spooler**.
7. Enter the full domain name of the network printer or its IP address. For example, printer3.myhost.example.com or 192.0.34.166.
8. On the **Configurable Parameters** list, enter 1 to choose item **1) Lp destination (queue) name** and assign the printer a unique local name (to be used in printing operations, such as 1p).  
Enter 0 to continue.
9. Enter y to cycle the print spooler and configure the new print queue.
10. Enter q, q, and press **Enter** to return to the **Printers and Plotters** screen.
11. Select **File**→**Exit** to return to HP SMH.

## Adding A Network Printer Using HP-UX Commands

To start JetDirect from the command line:

1. As a superuser, enter the command,  
`# /opt/hpnp/admin/jetadmin`
2. Enter 1 to choose option **1) Configuration (super-user only)**.
3. Enter 3 to choose option **3) Add printer to local spooler**.
4. Enter the full domain name of the network printer or its IP address. For example, `printer3.myhost.example.com` or `192.0.34.166`.
5. On the **Configurable Parameters** list, enter 1 to choose item **1) Lp destination (queue) name** and assign the printer a unique local name (to be used in printing operations, such as `lp`).  
Enter 0 to continue.
6. Enter `y` to cycle the print spooler and configure the new print queue.
7. Enter `q, q` to exit.

## Creating a Printer Class

A printer class is created when you assign the first printer to it with HP SMH or the `lpadmin` command. For information on what printer classes are, see *HP-UX System Administrator's Guide: Overview*.

You can assign a local printer or a network printer to a class.

You cannot assign a remote printer to a local printer class. However, you can create the class on the remote system and designate the class name as a remote printer on the local system.

You can use HP SMH to add a printer to a printer class when the printer is being added to the spooler; otherwise, you must use HP-UX commands.

## Creating a Printer Class Using HP-UX Commands

To use HP-UX commands, follow these steps after several printers have been added to the LP spooler:

1. Ensure that you have superuser capabilities.
2. Create the printer class, specifying a printer you want to add to the class of printers.

For example, to add a printer named `laser1` to the class of printers named `laser`, enter:

```
/usr/sbin/lpadmin -plaser1 -claser
```

Only one printer can be added to a class at a time. If you have more than one printer to add, repeat this command.

3. Allow print requests to be accepted for the newly added printer class. For example:

```
/usr/sbin/accept laser
```

4. Start the LP spooler if it is not already running:

```
/usr/sbin/lpsched
```

## Removing a Printer from the LP Spooler

These procedures can be used to remove local, remote, and network printers from the spooler.

### Removing a Printer Using Web-Based HP SMH

1. Ensure that an X Window System server is running on your local system. See “Using the X Window System” (page 32).
2. Start web-based HP SMH from your local system. See “Starting an HP SMH Web Client” (page 24).
3. Navigate to **Tools**→**Printer Management**→**Configure Printers or Plotters**.
4. Click the **Run** button.  
The **Printers and Plotters** window is displayed.
5. Highlight the printer or plotter you are removing.
6. Select **Actions**→**Remove**.



**NOTE:** HP SMH asks for confirmation before removing the printer from the LP spooler. If print jobs remain in the printer’s queue or if the printer is the system default destination, HP SMH notifies you. If you choose to remove a printer with jobs in its queue, HP SMH cancels them.

7. Click **OK** in the **Confirmation** dialog box.  
The printer queue is removed.
8. Select **File**→**Exit** to return to HP SMH.

### Removing a Printer Using HP-UX Commands

See *lpadmin(1M)* and *lpsched(1M)* for details on the command options.

1. Ensure that you have superuser capabilities.
2. (Optional) Notify users that you are removing the printer from the system.
3. Remove the printer from the configuration file of any software application through which the device is accessed. (Refer to the documentation accompanying the software application for instructions.)
4. (Optional) Deny any further print requests for the printer. For example:

```
/usr/sbin/reject -r"Use alternate printer." laser1
```

By doing this step, you can be assured that no new jobs will appear before you remove the printer.

Users will see the message `Use alternate printer` when they direct requests to a rejected destination if the printer has not been removed.

Once the printer has been removed and users try to send a request, they will see the message `Destination printer_name non-existent`.

5. (Optional) Determine if there are any jobs in the printer's queue. For example:

```
/usr/bin/lpstat -o laser1
```

6. (Optional) Disable the printer to be removed. For example:

```
/usr/bin/disable -r"Printer laser1 is disabled." laser1
```

You would issue the above `disable` command if there are jobs in the printer's queue and you do not want to wait for them to print before removing the printer. Issuing the `disable` command shuts the printer down in an orderly manner.

You can also specify the `-c` option on the `disable` command to cancel all print requests for the printer.

7. (Optional) If there are no jobs in the printer's queue, go on to Step 8. If there are jobs, decide whether to move all pending print requests in the request directory to another printer request directory or to cancel any requests. For example, to move print requests:

```
/usr/sbin/lpmove laser1 laser2
```

To cancel any requests:

```
/usr/bin/cancel laser1
```

8. Remove the printer from the LP spooler. For example:

```
/usr/sbin/lpadmin -x laser1
```

9. Start the LP spooler if it is not already running:

```
/usr/sbin/lpsched
```

## Removing a Printer from a Printer Class



---

**NOTE:** You cannot use HP SMH to remove a printer from a class.

---

### Removing a Printer from a Class Using HP-UX Commands

See `lpadmin(1M)` and `lpsched(1M)` for details on the command options.

1. Ensure that you have superuser capabilities.
2. Remove the printer from the class. For example, to remove printer `laser1` from class `laser`:

```
/usr/sbin/lpadmin -plaser1 -rlaser
```

3. Start the LP spooler if it is not already running:

```
/usr/sbin/lpsched
```

## Removing a Printer Class

---



**NOTE:** When you remove a printer class, the printers in the class are not removed — you may still use them as individual printers. If you remove all printers from a class, that printer class is automatically removed.

You cannot use HP SMH to remove a printer class.

---

### Removing a Printer Class Using HP-UX Commands

See *reject(1M)*, *lpmove(1M)*, *lpadmin(1M)*, and *lpsched(1M)* for details on the command options.

1. Ensure that you have superuser capabilities.
2. (Optional) Deny any further print requests for the printer class. For example:  
`# /usr/sbin/reject -r"Use alternate printer." laser`
3. (Optional): Determine if there are any jobs in the printer class's queue. For example:  
`# /usr/bin/lpstat -o laser`
4. (Optional) Move all pending print requests in the request directory for the printer class to another printer or printer class. For example:  
`# /usr/sbin/lpmove laser laser2`
5. Remove the printer class. For example:  
`# /usr/sbin/lpadmin -x laser`
6. Start the LP spooler if it is not already running:  
`# /usr/sbin/lpsched`

---

# 7 Configuring Mail

Whether you are administering a single system or a workgroup containing many systems, you will probably want your users to be able to communicate with each other using electronic mail (e-mail). This topic area will help you understand what is involved in setting up e-mail services for your workgroup.

## Components of an Electronic Mail System

To properly configure an electronic mail system, you need to know about the following components:

- “Mail User Agents” (page 121)
- “Mail Delivery Agents” (page 121)
- “Mail Alias Files” (page 122)
- “The Mail Queue” (page 122)
- “Networking Topographies” (page 123)
- “MIME Applications” (page 125)

### Mail User Agents

Mail User Agents are the programs that users run to send, and read e-mail. Mail User Agents that are shipped with HP-UX include `dtmail`, `elm`, `mail`, and `mailx`. Mozilla Thunderbird can be downloaded from the HP web site at <http://www.hp.com> . There are also commercially available Mail User Agents.

Although Mail User Agents appear to do all the work of transmitting and receiving e-mail, they are merely the visible part of the entire electronic mail system. Mail User Agents do not actually *deliver* the e-mail. Electronic mail *delivery* is handled by Mail Delivery Agents.

Mail User Agents:

- Format outgoing messages with proper header information and (if necessary) encode the outgoing messages for use by Mail Delivery Agents in routing the messages.
- Allow users to read, save, and delete incoming electronic mail messages.
- Schedule MIME Applications (if necessary) to allow the user to experience nontextual information attached to incoming electronic mail; for example, viewing graphics files or video clips, or listening to audio data.

### Mail Delivery Agents

Mail Delivery Agents form the core of the electronic mail system. These programs, usually running in the background, are responsible for routing, and delivering electronic

mail. On HP-UX and other UNIX systems, the primary Mail Delivery Agent is `sendmail`.

Although `sendmail` can be run directly from a shell command line to send a message, it is not usually used in this way. Mail User Agents are usually used as front ends to `sendmail` for sending mail.

### Mail Delivery Agents:

- Deliver mail to local users (users receiving e-mail on the computer that the Mail Delivery Agent is running on) by scheduling the `/usr/bin/mail` program or by forwarding the mail to users on local client machines.
- Forward e-mail via the appropriate transport mechanism not intended for local users to other computers/networks for delivery. For example, UUCP mail would be sent on its way by scheduling (and passing the message to) the `uux` program.
- Modify the format of the address information in message headers to accommodate the needs of the next computer or network in a message's delivery path, and to accommodate the delivery method that is being used to route the message. For example:

UUCP addresses are of the form:

`computername@domain.name!username`

whereas TCP/IP addresses can take one of several forms, such as:

`user`

`user@computer`

`user@computer.domain.name`

## Mail Alias Files

Mail Alias Files are used for:

- Mapping “real world” names to user login names
- Describing distribution lists (mailing lists), where a single name (for example, `deptXYZ`) is mapped to several or many user login names

For faster access, the alias files can be processed into a hashed database with the command `newalias` (a form of `sendmail`). By default, the system alias file (ASCII version) is located in the file `/etc/mail/aliases`.

## The Mail Queue

Outgoing messages cannot always be sent right away because of down computers, broken network connections, network traffic, and other reasons. Your Mail Delivery Agent needs a place to hold these messages until they can be sent on their way. That place is the mail queue.

If you are using `sendmail` (supplied with HP-UX) as your Mail Delivery Agent, your mail queue is, by default, the directory `/var/spool/mqueue`.

## Networking Topographies

Although there are many ways to configure electronic mail for a group of computers under your control, the following setups are often used:

- Central Mail Hub
- Gateway Mail Hub
- Fully Distributed

### Central Mail Hub

A central mail hub (a mail server) receives e-mail for its users and the users on the client computers that it serves. Users either NFS-mount their incoming mail files to their local computers (the clients), or log in to the hub to read their mail. Electronic mail can be *sent* directly from the client computers.

Advantages:

- Only one computer needs to be connected to the outside world, which protects (hides) the local clients from the network outside, giving the appearance that all mail from the workgroup is coming from a central computer.
- Only one computer needs to run the `sendmail` daemon (to “listen” for incoming e-mail).
- Data are centralized (easier to backup and control)

Disadvantages:

- Users of client machines must NFS-mount their incoming mail files from the hub (or log in to the hub) in order to read their mail.
- All electronic mail, *even between client machines in a local workgroup*, must go through the hub computer. This means that local mail traffic could be slowed if the hub machine becomes overloaded; and mail traffic would stop completely if the hub goes down or becomes disconnected from the network.

### Gateway Mail Hub

A gateway mail hub receives electronic mail for its users and users of client computers that it serves. The hub forwards mail intended for users of the client computers to those clients. Users do *not* NFS-mount their incoming mail files to their local (client) computers; they send and receive their mail directly from their own machines.

Advantages:

- Only one computer needs to be connected to the outside world, which protects (hides) the local clients from the network outside, giving the appearance that all mail from the workgroup is coming from a central computer.
- Traffic between local machines (within the workgroup) does not have to travel through the hub computer because

each client can send and receive its own electronic mail. Therefore if the hub goes down or becomes overloaded, local mail traffic is unaffected (only mail to and from computers outside of the workgroup is affected).

- Greater privacy for electronic mail users on the client machines. Data is not stored in a central repository.
- Each computer needs to run its own copy of the `sendmail` daemon to “listen” for incoming mail.
- Electronic mail from and to the outside world must travel through the hub, which could become a bottleneck if the mail traffic is heavy.

If the hub is down, clients cannot send and receive mail to and from computers outside of the work group.

#### Disadvantages:

### Fully Distributed

Each computer in the workgroup independently sends and receives its own electronic mail.

#### Advantages:

- There is no hub computer to contend with in this setup. Every computer, whether local to the workgroup or not, can send and receive electronic mail *directly* with every other computer in the network that also supports electronic mail.
- Greater privacy for electronic mail users on the individual machines. Data is not stored in a central repository.

#### Disadvantages:

- Because each computer (from an electronic mail perspective) is connected directly to the outside world, there is an increased data security risk.
- Each computer needs to run its own copy of the `sendmail` daemon to “listen” for incoming mail.

### Selecting a Topography

The topography you use depends on your needs. Here are some things to consider when choosing your electronic mail network topography:

#### Security

By using a topography with a hub computer you can better protect work that is being done on machines within your workgroup or organization. The single point of entry to your internal network (a gateway computer) is a lot easier to defend against unauthorized entry.

|                                        |                                                                                                                                                                                                                                                      |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Data Centralization                    | By having your mail files on a single machine or directory structure, it is easier to back up your data.                                                                                                                                             |
| Company Appearance and Future Planning | By using one of the topographies that use a hub computer, a small company can look more like a large corporation. As the company grows, the centralized mail processing can be easily moved to the jurisdiction of a corporate communications group. |
| Traffic Levels                         | If e-mail traffic levels are expected to be high, you might not want to use a single hub for processing all electronic mail.                                                                                                                         |

## MIME Applications

Gone are the days when electronic mail messages contained only ASCII text. Today people want to send other types of data: audio clips, still graphics (in a variety of formats), video clips, and so on.

Because Mail Delivery Agents were developed to handle the 7-bit ASCII data in text-only messages and not the 8-bit binary data contained in audio, graphics, and video, a method is needed for encoding the binary data to be transported by the text-only transport agents. The system developed for encoding the binary data is known as **MIME** (for Multipurpose Internet Mail Extensions).

Most modern Mail User Agents (including the CDE mail client, `dtmail` and the X-Window-enabled `elm`) can process MIME-encoded e-mail messages. For complete details about how MIME works, see RFC 1521. See also `elm(1)`.

## Configuring a System to Send Electronic Mail

To configure an HP-UX system to send e-mail, you need to do two things:

1. Be sure that the executable file for the Sendmail program, `/usr/sbin/sendmail`, is on your system.
2. If you are using a Gateway Mail Hub topography, you need to enable **site hiding** for each of the client computers in your workgroup, as described in “[Using Site Hiding](#)” (page 125).

## Using Site Hiding

With site hiding, the e-mail from users on client computers in your workgroup will appear to the outside world as if it were sent from the hub computer. Replies to such mail will be sent to the hub computer (unless a `Reply-To :` header in the e-mail directs otherwise).

1. On each *client* computer in the workgroup being served by a central mail hub, edit the file /etc/rc.config.d/mailservs:
  - a. Set the environment variable SENDMAIL\_SERVER to 0, indicating that this computer is not the hub, and is not a standalone e-mail system. The sendmail daemon will not be run on this computer:

```
SENDMAIL_SERVER=0
```
  - b. Set the environment variable SENDMAIL\_SERVER\_NAME to the official host name of the hub computer that will send and receive electronic mail on behalf of this client computer. For example, if the hub computer for a client has the official host name, corpmail.example.com, you would set the variable as follows:

```
SENDMAIL_SERVER_NAME="corpmail.example.com"
```
  - c. (Optional) The environment variable SENDMAIL\_FREEZE does not apply to clients, which always freeze the sendmail configuration file, but it is good practice to set this variable to 1 to indicate to viewers of the /etc/rc.config.d/mailservs file that the sendmail configuration file is being frozen for this client computer:

```
SENDMAIL_FREEZE=1
```
2. Reboot the client computer to enable site hiding and freeze the sendmail configuration file.

## Configuring a System to Receive Electronic Mail

Configuring a system in your workgroup to *receive* e-mail is a bit more complicated than configuring it to send e-mail. First you must determine two things:

1. Which type of networking topography you are going to use (see “Networking Topographies” (page 123))
2. Where the system fits in to the topography: the electronic mail hub, a client in a workgroup served by a hub, or a standalone system.

Using that information, begin by selecting the appropriate networking topography below:

- “Central Mail Hub Topography (Receiving E-mail)” (page 126)
- “Gateway Mail Hub Topography (Receiving E-mail)” (page 127)
- “Fully Distributed (Standalone System) Topography” (page 130)

### Central Mail Hub Topography (Receiving E-mail)

With this type of electronic mail system, a single computer serves as the place where all users in a workgroup send and receive e-mail. To do this, users either log in to the hub computer, or NFS mount their electronic mailboxes to local (client) systems. All

outgoing e-mail from the entire workgroup, even mail sent from a system that has NFS mounted an electronic mailbox, appears to have originated on the hub computer.

## Configuring the Hub

With Central Mail Hub topography, the electronic mail hub is the computer that receives e-mail from any computer outside of the workgroup on behalf of its own users and those of the client computers that it serves.

1. On the hub computer only, edit the file `/etc/rc.config.d/mailservs`:
  - a. Set the environment variable `SENDMAIL_SERVER` to 1 to indicate that this computer is the hub computer:  
`SENDMAIL_SERVER=1`
  - b. Set the environment variable `SENDMAIL_SERVER_NAME` to null to indicate that no other computer serves this one:  
`SENDMAIL_SERVER_NAME=`
  - c. (Optional) Set the environment variable `SENDMAIL_FREEZE` to 1 to indicate that the `sendmail` configuration file is to be frozen. With older computers, and in certain other circumstances, a frozen configuration file can speed up `sendmail`'s performance by reducing the time it needs to parse its configuration file.  
`SENDMAIL_FREEZE=1`
2. Reboot the hub computer to start up and properly configure the `sendmail` daemon.

## Configuring the Clients

With Central Mail Hub topography, the client computers do not receive electronic mail directly. Users either log into the hub computer to process electronic mail, or they NFS-mount their incoming mailbox files, typically located in the directory `/var/mount`, and run a Mail User Agent on their client system to process their mail. For outgoing mail (see “[Configuring a System to Send Electronic Mail](#)” (page 125)), the Mail User Agent will automatically schedule the `sendmail` program.

## Gateway Mail Topography (Receiving E-mail)

This type of electronic mail system is similar to the Central Mail Hub topography in that a single computer sends and receives e-mail on behalf of all of the users in the workgroup *to and from computers outside of the workgroup*. The difference is that e-mail within the workgroup e-mail does not have to go through the hub computer because each client machine is running its own copy of the `sendmail` daemon allowing it to receive e-mail directly from other computers in the workgroup.

## Configuring the Hub

The procedure for configuring the hub computer in a Gateway Mail Hub topography is:

1. On the hub computer, edit the file `/etc/rc.config.d/mailservs`:
  - a. Set the environment variable `SENDMAIL_SERVER` to 1 to indicate that this computer is the hub computer:  
`SENDMAIL_SERVER=1`
  - b. Set the environment variable `SENDMAIL_SERVER_NAME` to null to indicate that no other computer serves this one:  
`SENDMAIL_SERVER_NAME=`
  - c. (Optional) Set the environment variable `SENDMAIL_FREEZE` to 1 to indicate that the `sendmail` configuration file is to be frozen. With older computers, and in certain other circumstances, a frozen configuration file can speed up `sendmail`'s performance by reducing the time it needs to parse its configuration file.  
`SENDMAIL_FREEZE=1`
2. Reboot the computer to start up and properly configure the `sendmail` daemon.

## Configuring the Clients

Using Gateway Mail Hub topography each of the clients in a local workgroup can send e-mail to the others without having to go through the hub. For this to be successful each of the clients must be running its own `sendmail` daemon.

On each client computer:

- Edit the /etc/rc.config.d/mailservs file:
  - a. Set the SENDMAIL\_SERVER environment variable to 1. Although you are configuring a client computer in the workgroup, setting this environment variable to 1 will start the sendmail daemon each time you boot your client computer so that it can receive e-mail from other systems in your workgroup.  
SENDMAIL\_SERVER=1
  - b. Set the SENDMAIL\_SERVER\_NAME environment variable to the name of the computer that will be the gateway to the outside world. For example, if the gateway computer was called gateway.corp.com:  
SENDMAIL\_SERVER\_NAME="gateway.corp.com"
  - c. (Optional) The environment variable SENDMAIL\_FREEZE does not apply to clients (which always freeze the sendmail configuration file), but it is probably good practice to set this variable to 1 to indicate to viewers of the /etc/rc.config.d/mailservs file that the sendmail configuration file is being frozen for this client computer:  
SENDMAIL\_FREEZE=1

## Fully Distributed (Standalone System) Topography

When using a Fully Distributed electronic mail topography, each computer is a standalone machine (with regard to electronic mail). Each machine is effectively its own workgroup and is configured just like the hub computer in a Central Mail Hub topography e-mail network.

### Configuring Each System

The procedure for configuring each system in a Fully Distributed topography is:

1. Edit the file /etc/rc.config.d/mailservs:
  - a. Set the environment variable SENDMAIL\_SERVER to 1 to indicate that this computer will run the sendmail daemon to receive mail:  
`SENDMAIL_SERVER=1`
  - b. Set the environment variable SENDMAIL\_SERVER\_NAME to null to indicate that no other computer serves this one:  
`SENDMAIL_SERVER_NAME=`
  - c. (Optional) Set the environment variable SENDMAIL\_FREEZE to 1 to indicate that the sendmail configuration file is to be frozen. With older computers, and in certain other circumstances, a frozen configuration file can speed up sendmail's performance by reducing the time it needs to parse its configuration file.  
`SENDMAIL_FREEZE=1`
2. Reboot the computer to start up and properly configure the sendmail daemon.

---

## 8 Setting Up the Online Manpages

There are three ways to set up online manpages, each resulting in a different amount of disk usage and having a different response time:

1. Fastest response to the `man` command (but heaviest disk space usage):

Create a formatted version of *all* the manpages. This is a good method if you have enough disk space to hold the nroff originals and the formatted pages for the time it takes to finish formatting. To start the formatting process, enter:

```
catman
```

Formatting all the manpages can take some time, so you might want to run the process at a lower priority.

2. Medium response time to the `man` command (with medium disk space usage):

Format only heavily used sections of the manpages. To format selected sections, enter:

```
catman sections
```

where *sections* is one or more logical section IDs from the *HP-UX Reference*, such as 1, 1m, 2, 3. Note that the letter m is lowercase. For example,

```
catman 1m357
```

3. Slowest response to the `man` command (but lightest disk space usage):

Do not format any manpages. HP-UX will format each manpage the first time a user specifies the `man` command to call up a page. The formatted version is used in subsequent accesses (only if it is newer than the unformatted source file).

To improve response time, you can make directories to hold the formatted manpages. To determine the directory names you need, check the `MANPATH` variable. For example, to create directories for the default `/usr/share/man` directory, execute the following script:

```
cd /usr/share/man
```

```
mkdir cat1.Z cat1m.Z cat2.Z cat3.Z cat4.Z cat5.Z \
> cat6.Z cat7.Z cat8.Z cat9.Z
```

You only need to create the `cat8.Z` directory if `/usr/share/man/man8.Z` exists. To save disk space, make sure you use the `cat* .Z` directories (not `cat* .Z`) because if both `cat* .Z` and `cat* .Z` exist, both directories are updated by `man`.

To save disk space, you can NFS mount the manpages on a remote system.

Regardless of how you set up the manpages, you can recover disk space by removing the nroff source files.



**CAUTION:** Before removing any files, make a backup of the man directories you created in case you need to restore any files.

For example, to remove files for section 1 in /usr/share/man, enter:

```
rm man1/*
rm man1.Z/*
```

This concept for recovering disk space also applies to localized manpages. For further details, see *man(1)* and *catman(1M)*.

---

# 9 Configuring Peripherals

To add peripherals to your system, consult the following documentation:

- The hardware installation document that came with the peripheral.
- For PCI OL\* information, see the *Interface Card OL\* Support Guide*. For PCI OL\* information on nPartition-able systems, see the *nPartition Administrator's Guide*.

PCI OL\*, previously known as OLAR, is the ability to add or remove a PCI card without needing to completely shutdown the entire system. The system hardware combined with operating system support allows per-slot power control. Instead of turning off the entire system, you can turn off and on power to a specific PCI slot.

The procedures for PCI OL\* can be performed through the Peripheral Devices section of HP SMH, pdweb, or the Partition Manager for nPartition-able systems, or through HP-UX commands, such as `olrad`. All are documented in the preceding documents.



**CAUTION:** Before attempting these procedures, please read the documents mentioned above. Turning off power to certain PCI slots can have disastrous effects; for example, if the PCI slot connects to an unmirrored root or swap disk, the system will crash. Further, the I/O card itself needs to be checked for OL\* functional compatibility as well as compatibility to the specific PCI slot; for example, you cannot insert a 33 MHz card in a slot running a 66 MHz bus.

- For general peripherals, see *Interface Card OL\* Support Guide*.
- See the *HP-UX 11i Release Notes* for the titles of documents that may be relevant to installing peripherals. Such documents may contain specific information on the software driver and the device special file for communication with particular peripherals.

For HP-UX to communicate with a new peripheral device, you may need to reconfigure your system's kernel to add a new driver. See “Configuring the Kernel” (page 153) for details.

## Configuring PCI Error Recovery

The PCI Error Recovery feature provides the ability to detect, isolate, and *automatically* recover from a PCI error, avoiding a system crash. It is included with the HP-UX 11i v3 operating system and is *enabled* by default.

To enable and disable PCI Error Recovery, see “Controlling PCI Error Recovery” (page 134).

## What is PCI Error Recovery?

If PCI Error Recovery is enabled and an error occurs on a PCI bus containing an I/O card that supports PCI Error Recovery, the following steps are taken:

1. The PCI bus is quarantined to isolate the system from further I/O and prevent the error from damaging the system.
2. The PCI Error Recovery feature attempts to recover from the error and re-initialize the bus so I/O can resume.

If an error occurs during the automated error recovery process, the bus and I/O card will remain quiesced.

If the bus contains a card that supports online addition, replacement, or deletion (OL\*) and the card is in a hot pluggable slot, you can use the `olrad` command (or the attention button) to manually recover from the error by replacing the card.

For information on OL\* operations, see the *Interface Card OL\* Support Guide*. To determine if OL\* is supported, see the documentation or support matrix for the specific I/O card.

If the PCI Error Recovery feature is *disabled* and an error occurs on a PCI bus, a Machine Check Abort (MCA) or a High Priority Machine Check (HPMC) will occur and the system will crash.



**CAUTION:** If you use HP Serviceguard, HP recommends that you enable the PCI Error Recovery feature only if your storage devices are configured with multiple paths and you have *not* disabled HP-UX native multipathing. If PCI Error Recovery is enabled, but your storage devices are configured with only a single path, HP Serviceguard may not detect when connectivity is lost. HP Serviceguard will not cause a failover unless it detects a loss of connectivity.

## Controlling PCI Error Recovery

PCI Error Recovery is controlled by two tunables that you can configure, using HP SMH, kcweb, or kctune. See “Managing Kernel Tunable Parameters with `kctune`” (page 170) and “Managing Kernel Tunable Parameters with HP SMH” (page 175).

- `pci_eh_enable`

This tunable enables or disables the PCI Error Recovery feature. It is enabled by default. Since `pci_eh_enable` is not a dynamic tunable, a reboot is required for changes to take effect.

- `pci_error_tolerance_time`

This tunable determines whether an automatic PCI error recovery will occur on an I/O slot, based on the time interval between two PCI errors. If two PCI errors occur on a PCI slot within the time interval specified by

`pci_error_tolerance_time`, the card in the I/O slot will be suspended and you will need to attempt a manual recovery operation to restore the card.

## PCI Error Recovery Documentation

PCI Error Recovery is supported by the following documentation, available on the HP Technical Documentation web site at <http://docs.hp.com>:

In the High Availability section:

- *PCI Error Recovery Product Note*
- *PCI Error Recovery Support Matrix*
- *Interface Card OL\* Support Guide*

In the *HP-UX Reference*:

- `pci_eh_enable(5)` manpage
- `pci_error_tolerance_time(5)` manpage

## Ethernet Configuration and Verification

See also the *HP-UX LAN Administrator's Guide*.

Use this procedure to configure and verify any type of HP-UX 11i v3-based Ethernet type cards including 10 Gigabit Ethernet, Gigabit Ethernet, and Fast Ethernet.

### Procedure 9-1 Installation, Configuration, and Verification Procedure

1. Install LAN card (if needed) and attach cables to it according to instructions in the server's hardware support document.

2. Plug LAN card into the appropriate switch port.

3. Boot the system.

4. To see all LAN cards with drivers, run:

```
ioscan -fnC lan
```

5. If you do not see your card, use `ioscan` to determine if the card is there but needs a driver (UNCLAIMED).

```
ioscan -fn
```

If your LAN card is a built-in card or was factory installed, the software bundle required for the card is already loaded onto your system's hard drive.

6. Install or update the driver if needed (to the latest version). Get the latest driver from the latest quarterly update media for your version of HP-UX. To add new software, see "Installing or Updating a Networking Driver with `swinstall`" (page 136).

7. Verify connectivity to the switch.

```
nwmgr -c lanPPA -A link_state,speed
```

8. View cards with IP addresses.

- ```
# netstat -in
```
9. Assign IP addresses and subnet masks to additional cards as desired (options given later). See “Assigning IP Addresses and Subnet Masks to Additional Cards Using Web-Based HP SMH” (page 137).
 10. For each interface, run ifconfig to verify the correct IP, netmask, and broadcast addresses. For example:


```
# ifconfig lan0
```
 11. For each interface, ping the broadcast address shown in Step 10 and verify that you get responses from other nodes on that subnet.

This should get the system up and connected to any local networks. To set up communication between the local network and remote networks, see “Routing: Configuring the LAN to Reach Other Networks” (page 139).

Installing or Updating a Networking Driver with swinstall

You can either load the entire HP-UX 11i v3 operating environment (OE) from the distribution media and you will automatically get the correct LAN (and if applicable, mass storage) software bundles, or you can select and load the software bundle (or bundles) required for your card. The HP-UX Ethernet Card System-Driver Matrix, located in the document *Ethernet System-Driver Matrix for HP-UX 11i v1, 11i v2, and 11i v3* in the I/O Cards and Networking Software section of <http://docs.hp.com>, tells which drivers are used with each card. To load a specific bundle:

1. Load the software media into the appropriate drive.
2. Run the swinstall command. See *swinstall(1M)* for details.
3. In the **Select Source** dialog, identify the depot containing the software by setting the **Source Host Name** and **Source Depot Path**, as necessary. Click **OK**.
4. In the **Software Selection** window, highlight the desired software bundle, such as the 1000Base-T software driver bundle GigEther-01 (for cards such as A6825A/A6847A 1-port) or IEther-00 (for cards such as A7011A/A7012A 2-port and AB545A 4-port).
5. Choose **Mark for Install** from the **Actions** menu to choose the bundle.
6. Choose **Install** from the **Actions** menu to begin product installation and open the **Install Analysis** window.
7. Click **OK** when the **Status** field displays Ready.
8. Click **Yes** in the **Confirmation** window to confirm that you want to install the software. swinstall loads the bundle, runs the control scripts, and builds the kernel. This should take about 3 to 5 minutes. When the **Status** field indicates Ready, click **Done**.
9. A **Note** window opens. Click the **OK** button to reboot the system.

Assigning IP Addresses and Subnet Masks to Additional Cards Using Web-Based HP SMH

To configure the LAN cards, you can use the web-based HP System Management Homepage (HP SMH).

1. Ensure that an X Window System server is running on your local system. See “Using the X Window System” (page 32).
2. Start web-based HP SMH from your local system. See “Starting an HP SMH Web Client” (page 24).
3. Navigate to **Tools**→**Network Interfaces Configuration**→**Network Interface Cards** (the **HP-UX Network Interfaces Configuration Tool** screen) and select the **NIC** tab.
4. Select the radio button of the LAN card that you want to modify. The interface details are displayed below the list.
5. In the procedure list on the right side, click **View/Modify IP Attributes**. The **View/Modify IP Attributes** screen is displayed.
6. Configure the IP address, host name, and any other desired parameters. Click the question mark (?) in the upper right corner of the screen to consult the online help.
7. (Optional) Select **Preview** to see the commands that will be executed.
8. Click **OK** to continue or **Cancel** to quit the process.

Frequently Asked Questions

This section describes how to solve some of the most frequently encountered problems and answers the most frequently asked questions.

1. Driver information: What driver do I need ?
 - a. What is the driver name in HP SMH (or other tools), for example, iether? See the Ethernet Support matrixes located in the document *Ethernet System-Driver Matrix for HP-UX 11i v1, 11i v2, and 11i v3* in the **I/O Cards and Networking Software** section of <http://docs.hp.com>.
 - b. Is the software on the system?The principal Gigabit Ethernet and Fast Ethernet drivers are preloaded onto your system's hard drive; the 10 Gigabit Ethernet driver needs to be selected and loaded from the OE media. To see if the driver software bundle is loaded on your hard drive, run swlist.
 - c. How can I tell if the driver is in the kernel?See “Verifying that the Networking Driver is in the Kernel” (page 142).
 - d. What software do I load?Most networking drivers are always installed, so you don't have to load them onto your hard drive. The ones that are preloaded need to be configured to

add IP addresses, and if desired, subnet masks, and any other optional settings. The 10 Gigabit Ethernet driver `ixgbe` is not pre-loaded; it is *selectable*. So, when setting up 10 Gigabit Ethernet networking for the first time on a system, you need to load it onto your hard drive using the `swinstall` utility.

e. Where can I get the software?

All of the currently used networking drivers are included in the latest quarterly HP-UX update. The drivers are also available on the web. Go to <http://www.hp.com> and then choose **Software and Driver Downloads**.

f. How can I tell which version I have installed?

See “Verifying that the Networking Driver is in the Kernel” (page 142).

2. Am I connected to the switch and at what speed?

In most cases, this a speed/duplex question. For details, see “Setting Speed and Duplex Mode of the LAN Card and Link Partner (Switch or Router)” (page 140).

Verifying LAN Installation

1. Verify that the LAN connector’s Link LED is steadily on. This means the card and driver are installed successfully.
2. Obtain the card name and the station address of each card by using the `nwmgr` command. The MAC address labeled on each card refers to LAN port A (the right port). Add 1 for each additional port to obtain the MAC address for any additional LAN ports.
3. To verify link-level connectivity with a remote system, enter:

```
# nwmgr --diag -c cardname -A dest=0x00306E2DF7FE
```

For example,

```
# nwmgr --diag -c lan0 -A dest=0x00306E2DF7FE
```

When you use `nwmgr`, ensure that the remote system is on the same LAN segment and is an HP-UX-based system.

4. To verify IP-level connectivity with a remote system, enter:

```
# ping remote-address -n 5  
# netstat -in
```

Each time you run the command pair, the values for `Ipkts` and `Opkts` should increase.

Installation is complete when you have successfully run `nwmgr`, `ping` and `netstat`.

Optionally, if you want to verify that the driver appears for each installed card, enter:

```
# ioscan -fknClan
```

The output for each port would look something like the following:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
lan	6	1/0/2/1/0/6/0	iether	CLAIMED	INTERFACE	HP AD193A PCI/PCI-X 2-port 4Gb FC/2-port 1000B-T Combo Adapter
lan	6	1/0/2/1/0/6/1	iether	CLAIMED	INTERFACE	HP AD193A PCI/PCI-X 2-port 4Gb FC/2-port 1000B-T Combo Adapter

The last two digits of the hardware path (third column) reflect the path of each port; in the sample output shown, the 0 indicates LAN A and the 1 indicates LAN B (for the two LAN ports on a multiport card). In this example of a card with two LAN ports, both ports need to show as CLAIMED here.

Routing: Configuring the LAN to Reach Other Networks

Once your IP address and subnet mask have been set, HP-UX should be able to deliver data to any node on your local network. In order for your local network to reach other networks, your machine needs access to a router or switch (devices that route data to other networks). A router is also considered a gateway to another network. Configuring the identity of a default gateway on your local machine can be done in any of four ways:

- Using HP SMH.
- Using `set_parms`.
- Editing the `/etc/rc.config.d/netconf` file. The following entries define the routing for a single interface on a system.

```
ROUTE_DESTINATION[0] = "default"
ROUTE_MASK[0] = ""
ROUTE_GATEWAY[0] = "196.6.20.2"
ROUTE_COUNT[0] = "1"
ROUTE_ARGS[0] = ""
```

As needed, set `ROUTE_DESTINATION[0] = "default"` and `ROUTE_COUNT[0] = "1"`. Save changes.

After editing the `netconf` file, to start services and initiate the new route, use:

```
# /sbin/init.d/inetd start
# /sbin/init.d/net start
```

- Using the `route` command to put the new route into effect on the system as follows:

```
# /usr/sbin/route add default router_ip_address 1
```

If a different default route has already been defined, avoid ending up with two default routes on the system—one to the old gateway and one to the new gateway—by removing an existing route first. Use:

```
# route -f
# route add default router_ipaddress 1
```

If your system does not need to access another system on a different network, you can assign the local system's IP address as the `ROUTE_GATEWAY`, because it will be routing all network traffic directly to other local area systems. If the local system will be

communicating with systems outside the local network, the ROUTE_GATEWAY must be the IP address of the gateway router. The system maintains a dynamic routing table in memory. This routing table is dynamic information on how to route to a specific system or network. You can add and delete routes in this table by using the route command. You can force a specific path to a destination if there is more than one way to get there. For details, please see the *route(1M)* manpage.

To verify the configuration, use the netstat command to display the current routing table.

```
# netstat -rn
Routing tables
Destination        Gateway            Flags  Refs Interface   Pmtu
127.0.0.1          127.0.0.1         UH    0      lo0          4608
12.56.219.151     12.56.219.151   UH    59     lan0         4608
12.56.216.0        12.56.219.151   U     7      lan0         1500
127.0.0.0          127.0.0.1         U     0      lo0          0
default            15.0.64.1        UG    0      lan0         0
```

Setting Speed and Duplex Mode of the LAN Card and Link Partner (Switch or Router)

Gigabit Ethernet Base-T connections on those cards operate at 10 or 100 megabit/s in either full- or half-duplex modes and at 1000 megabit/s only in full-duplex mode.

When the LAN autonegotiates, it achieves the highest performance if the link partner (switch or router) is also set to either autonegotiation or 1000FD. The following table shows the resulting speed depending on the settings of the ports on the card and on your router or switch.

Table 9-1 HP-UX 1000Base-T Supported Configurations

HP-UX 1000Base-T Port	Link Partner	Resulting Speed
AUTO	AUTO	Highest Common Speed (HP-UX supports 10/100/1000)
AUTO	1000 FD fixed/manual	1000 Mb/s FD
10 HD	10 HD (for example, a 10Base-T Hub)	10 Mb/s HD
10 FD	10 FD	10 Mb/s FD
100 HD	100 HD	100 Mb/s HD
100 FD	100 FD	100 Mb/s FD

Ensuring Card and Link Partner Speed and Duplex Settings

A lot of duplex mismatch issues can show up as other problems. The following nwmgr results can tell whether the card and link partner are operating as you need.

1. Obtain the PPA number of each LAN link you are testing with nwmgr.

2. Test the card and link with nwmgr, as follows:

```
# nwmgr -c lanPPA -A link_state,speed
```

For example,

```
# nwmgr -c lan0 -A link_state,speed
```

lan0 current values:

Link State = Up

Speed = 1 Gbps Full Duplex (Autonegotiation : On)

The Speed line includes the current speed, the current duplex setting, and how the setting was determined.

If the last item says Autonegotiation : On (as above), this indicates you have set the card to auto_on and the card determines the correct setting. On Fast Ethernet links, this is only appropriate if the switch port is also in auto_negotiate mode. It will not work properly if the switch is hard set to 100FD. If you see something like:

Speed = 100 Mbps Half Duplex (Autonegotiation : On)

(note the Half Duplex) then it almost certainly means that the switch is hard set to 100FD. In this case, you will encounter collisions and dropped packets and potentially poor network performance or even see the switch disable the network port due to a (perceived) high error rate.

If you see:

Speed = 100 Mbps Full Duplex

(note the missing (Autonegotiation : On)), it means that your card is hard set to 100FD and is not autonegotiating. In this case, the switch must also be hard set or you will see errors as noted above.

HP recommends the settings shown in [Table 9-1 \(page 140\)](#). For 100Mb/s links, setting one side of the link to autonegotiate and the other hard set to 100FD will cause connectivity problems. Always verify that the switch setting is the recommended one used across your site (such as 100FD).

Configuring Optional Jumbo Frames Size for Gigabit Ethernet

(Jumbo frames are supported only at 1000 Mb/s)

- Jumbo frames for the iether and igelan drivers have a maximum transmission unit (MTU) in the range 1501 to 9000 bytes. Normal frames on have an MTU in the range 256 to 1500. If you are using Jumbo Ethernet frames, ensure that:
 - All end stations on a given LAN have the same MTU setting;
(In the Jumbo Frames description, “LAN” means that the end stations do not have any routers or layer 3 switches between them.)
 - Intermediate stations such as switch ports in your LAN have an MTU equal to or greater than the end station’s MTU.
- Obtain the PPA number of the card by entering nwmgr.

- Start web-based HP SMH as described in “Starting an HP SMH Web Client” (page 24). Then:
 1. Navigate to **Tools**→**Network Interfaces Configuration**→**Network Interface Cards**.
 2. Select the **NIC** tab.
 3. Select the LAN interface by clicking the radio button.
The interface details are displayed below the list box.
 4. In the right hand actions column, click **View/Modify NIC Attributes**.
The **View/Modify NIC Attributes** window is displayed.
 5. To use Jumbo Frames with the *iether* and *igelan* drivers, set the MTU value to a number in the range 1501 to 9000 bytes.
The interface will be configured for Jumbo Frame operation, which will be retained across reboots.
- Verify the MTU change by entering `netstat -rn`. If the MTU has not changed, enter the following commands (as root):


```
# ifconfig lanPPA unplumb
# ifconfig lanPPA ip_address netmask netmask up
```

Verifying that the Networking Driver is in the Kernel

The Gigabit Ethernet drivers updated in December 2006 work on HP-UX 11i v3. To verify that the driver was loaded in the kernel, execute the following command:

```
# what /stand/vmunix |grep drivername
```

where *drivername* is *igelan* for the GigEther-01 bundled Gigabit Ethernet or *iether* for the IEther-00 bundle.

You should see a response like:

```
igelan_ilan Version: 1 Dec 18 2006
igelan Revision: B.11.31.05 Dec 18 2006
```

Related System and Network Administration

Refer to Table 9-2 (page 143) to determine what configuration tool is best for a specific job and where to find information on them.

Table 9-2 Choosing a Configuration Tool

Associated system and networking tasks	Use tool or command	For further information
If setting up first network access on a system, you will need an IP address, and a host name.	<code>set_parms</code> You can configure the first or "core" LAN even if the configuration was postponed until after startup.	See "Setting System and Network Parameters" (page 41) or <i>set_parms(1M)</i> .
If adding multiple LAN cards:	To assign static IP addresses, use HP SMH.	See "Assigning IP Addresses and Subnet Masks to Additional Cards Using Web-Based HP SMH" (page 137).
	To assign dynamic IP addresses, use DHCP.	See the <i>HP-UX IP Address and Client Management Administrator's Guide</i> or the <i>Ignite-UX Administration Guide</i> .
	To resolve host names, use the domain name service (DNS).	See the <i>HP-UX IP Address and Client Management Administrator's Guide</i> .
	To aggregate ports for improved load balancing and failover, use HP Auto Port Aggregation (APA).	See the <i>HP APA Support Guide</i> .
If configuring multiple systems, and installing multiple OE images:	Use Ignite-UX.	See the <i>Ignite-UX Administration Guide</i> .
If setting up network services:	Enable Internet Services governed by <code>inetd</code> . Automatically done by start-up script.	See <i>Using Internet Services</i> .

Further Ethernet Information

For further information on all the current 10 Gigabit, Gigabit, and Fast Ethernet, see the I/O Cards and Networking Software section on the HP Technical Documentation web site, <http://docs.hp.com>.

For maintenance and troubleshooting information about the current 10 Gigabit, Gigabit, and Fast Ethernet drivers, see the *Ethernet Support Guide*

Configuring Non-HP Terminals

To set up a user with a non-HP terminal, do the following:

1. Make sure the file set NONHPTERM is on the system by using either of these methods:

- **# swlist -l filesset NonHP-Terminfo**
If the file set exists, the entry for NonHP-Terminfo.NONHPTERM will be displayed.
- **# ll /var/adm/sw/products/NonHP-Terminfo**
If the file set exists, the directory /var/adm/sw/products/NonHP-Terminfo/NONHPTERM will exist.

If the fileset is not on the system, you will need to load it from your latest HP-UX media. For details, see the *HP-UX System Administrator's Guide: Routine Tasks* or the *Software Distributor Administration Guide*.

2. Look in the directory /usr/share/lib/terminfo for a file that corresponds to the terminal you want to set up. For example, suppose you want to set up a user with a Wyse 100 terminal. All supported terminals whose names begin with w are contained in the /usr/share/lib/terminfo/w directory. Because this directory contains an entry wy100, you have probably found the correct file. To be sure, examine the contents of the file with more. You will see a screen full of special characters, but near the beginning you will see wy100|100|wyse 100. This verifies the correct file and shows that you can refer to the Wyse 100 by any of the names wy100, 100, or wyse 100.

If there is a terminfo file for the terminal you want to add, skip Step 3 and go to Step 4.

If there is no terminfo file for the terminal you want to add, you will need to create one. See Step 3 for details.

3. To create a terminfo file, follow the directions in *terminfo(4)*.

To adapt an existing file, follow these steps:

- a. Log in as superuser.
- b. Make an ASCII copy of an existing terminfo file. For example, make a copy of the file /usr/share/lib/terminfo/w/wy100 by entering:
untic /usr/share/lib/terminfo/w/wy100 > new_file
- c. Edit the new file to reflect the capabilities of the new terminal. Make sure you change the name(s) of the terminal in the first line.
- d. Compile the new terminfo file:
tic new_file

For more further information, see *tic(1M)* and *untic(1M)*

4. Set the user's TERM variable in the appropriate login script (either `.profile` for Korn and POSIX shell users or `.login` for C shell users in their home directory) to any of the names you uncovered in Step 2. For example:

```
# export TERM=wy100 (Korn or POSIX shell)
# setenv TERM wy100 (C shell)
```

The default versions of these scripts prompt the user for the terminal type upon log in, so rather than editing the script, you could simply tell the user to respond with the terminal name. For example:

```
TERM = (hp) wy100
```

You can also set the TERM variable with the `/sbin/ttypype` command.

Troubleshooting Terminals

There are a number of terminal-related problems that can occur. Many of these result in a terminal that appears not to communicate with the computer. Other problems cause "garbage" to appear on the screen (either instead of the data you expected or intermixed with your data).

This section primarily addresses problems with alpha-numeric display terminals; however, many of the steps discussed here can also be applied to problems with terminal emulators such as HP AdvanceLink (running on a Vectra PC) or X Window terminal processes (such as `hpterm` and `xterm`). Also see "[Other Terminal Problems](#)" (page 151).

Unresponsive Terminals

There are many things that can cause a terminal not to respond (no characters are displayed except, perhaps, those which are displayed by the terminal's local echo setting). Here is a procedure you can use to find many of them.

1. Determine the status of the system.

Is the system still up? If not, you've probably found your problem. You will need to reboot the system.

Is the system in single user state? If so, the only active terminal will be the system console. Other terminals will not respond. You will need to switch to a multiuser state. See the *init*(1M) manpage for more information on changing run states.



NOTE: To determine the run state of your system (from a working terminal), enter:

```
$ who -r
```

The output will look something like:

```
.      run-level 3  Jun  3 22:25      3          0      S
```

The current state of the machine is in the field immediately to the right of the time (third field from the right). For complete information on each of the fields, consult the *who*(1) manpage.

2. Determine if an editor is running on the terminal.

This is best done from another terminal. Issue the command:

```
$ ps -ft terminal
```

This displays *all* processes associated with the *terminal* with which you are having problems. For each entry, check in the column marked **COMMAND** to see if the process represented by that entry is an editor.

If you find that an editor *is* running at the terminal, it is probably in a text-entry mode. You will need to save the work and exit the editor. For directions on how to do this, consult the manpage for the appropriate editor.



CAUTION: If you are not sure of the status of the work being edited, *DO NOT* simply save the file and exit. You will overwrite the previous contents of the file with unknown text. Save the work in progress to a temporary file so that both the original and edited versions of the file are accessible.

3. Enter **Ctrl-Q** at the terminal keyboard.

Terminals frequently use the start/stop (XON/XOFF) protocol to start and stop output to them. **stop** is usually defined as **Ctrl-S** (XOFF) and **start** as **Ctrl-Q** (XON). If output to the terminal was stopped because a **stop** signal was sent from the terminal to the computer, it can be restarted by sending the computer a **start** signal (for example, type **Ctrl-Q** from the problem terminal's keyboard). Sending the **start** signal does not harm anything even if no **stop** signal was previously sent.

If the problem is an application program that's looping or not functioning properly, try pressing the **Break** key and then try the `intr` signal (usually **Ctrl-C**) or the `quit` signal (usually **Ctrl-l**) to see if you can get a shell prompt back. To find out what the `intr` or `quit` signal is for the affected terminal, go to a working terminal and enter the command:

```
# stty -a < /dev/terminal-device
```

For example:

```
# stty -a < /dev/ttyp1
```



NOTE: The `stty` command, above, should only be used with device file names for *currently active* terminal device files (use the `who -R` command to see which device files are active). If you attempt to execute `stty` with a nonactive tty device file, the command may hang, waiting for input. Press **Ctrl-C** to abort it.

4. Reset the terminal.

The terminal itself may be stuck in an unusable state. Try resetting it. Consult your terminal owners document for information on how to do this. Powering the terminal off, waiting for a few seconds and powering it back on will also reset the terminal.



NOTE: Power cycling a terminal can have the same effect as sending a BREAK, which can make the host think it got a BREAK and change the baud rate. If this happens a lot, use a `gettydefs` entry that does not cycle through baud rates.

5. Check the terminal configuration.

The terminal may not be configured correctly. You should verify the following:

- Is the terminal in Remote * mode? *It should be.*
- Is Block * mode turned ON? *It should not be.*
- Is Line * mode turned ON? *It should not be.*
- Is Modify * mode turned ON? *It should not be.*

6. Check the physical connection.

Check to make sure that:

- All cables are firmly attached and in their proper locations.
- All interface cards are firmly seated in their slots.
- The power cord to the terminal is firmly connected.
- The power switch is turned on.

7. Kill processes associated with the problem terminal.



CAUTION: Use *extreme caution* when killing processes. The processes will be immediately and unconditionally terminated. Some valid processes might take a long time to complete. Be sure to type carefully when entering the PID numbers for the `kill` command to avoid killing the wrong process.

If you have another terminal that is still working, go to that terminal and log in (you will need to be superuser). Execute the command:

```
# ps -ft terminal
```

This displays *all* processes associated with the `terminal` with which you are having problems. Look at the column marked PID (these are the process IDs for the processes associated with that terminal). Execute the following command, listing each process ID associated with the problem terminal:

```
kill -9 process-id
```

For example:

```
# kill -9 20133
```

This should kill all processes associated with that terminal. The `init` process will then respawn a `getty` process for that terminal (if it has been set up to do that, in the `/etc/inittab` file) and you should once again be able to log in.

8. Attempt to log in to the previously hung terminal again.

If you are successful, you've fixed the problem. If not, continue to the next step.

9. Use `cat` to send an ASCII file to the hung terminal's device file.

HP-UX communicates with peripherals through device files. These special files are typically located in the directory `/dev` and are used by HP-UX to determine which driver should be used to talk to the device (by referencing the **major number**) and to determine the address and certain characteristics of the device with which HP-UX is communicating (by referencing the **minor number**).

Try using the `cat` command to send an ASCII file (such as `/etc/motd` or `/etc/issue`) to the device file associated with the problem terminal. For example, if your problem terminal is associated with the device file `ttyd1p4`:

```
# cat /etc/motd > /dev/ttyd1p4
```

You should expect to see the contents of the file `/etc/motd` displayed on the terminal associated with the device file `/dev/ttyd1p4`. If you do not, continue to the next step.

10. Check the parameters of the device file for the problem terminal.

Device files have access permissions associated with them, just as other files do. The file's access permissions must be set so that you have access to the file. If you set the file's permissions mode to 0622 (`crw--w--w-`), you should be safe.

If the file's permissions are set to allow write access and the file isn't displayed on the terminal, check the major and minor numbers of the device file. You can list them with the `ll` command. You can use the `ls -f` command to interpret the major and minor numbers and display the results.

11. Other things to check.

- Make sure your `inittab` entries are active

If you are just adding this terminal and have made a new entry in the `/etc/inittab` file by editing it, remember that this doesn't automatically make your new entry active. To do that you need to enter the command:

```
# init -q
```

This tells the `init` process to scan the `/etc/inittab` file to update the information in its internal tables.

- Check for functioning hardware.

Now is the time to check the hardware. To do this, check the following items:

- If your terminal has a self-test feature, activate it. If not, power the terminal off, wait several seconds, and power the terminal back on. This will test (at least to some degree) your terminal hardware.



NOTE: Power cycling a terminal can have the same effect as sending a BREAK, which can make the host think it got a BREAK and change the baud rate. If this happens a lot, use a `gettydefs` entry that does not cycle through baud rates.

- An alternate method to test the terminal hardware is to swap the suspect terminal with a known good one. This will help identify problems within the terminal that are *not* caught by the terminal selftest.



NOTE: Be sure to swap only the terminal (along with its keyboard). You want the known good terminal at the end of the SAME cable that the suspect terminal was plugged into. Also, plug the suspect terminal (with its keyboard) into the same cable that the known good terminal was plugged into and see if it functions there.

- If the known good terminal doesn't function on the suspect terminal's cable, and the suspect terminal is working fine in its new location, you can be confident that the terminal itself is functioning properly and the problem is elsewhere.
- The next thing to check is the cable connecting the terminal to the computer. Swap the suspect cable with a known good one.



NOTE: Since you know the terminal at the end of each cable is working, you only have to swap the ends of the cables where they connect to the computer. If the problem remains with the terminal it was associated with prior to the cable swap, you probably have a broken or miswired cable. If the problem transfers to the other terminal (and the previously bad

terminal/cable combination works in its new location), then the problem is most likely with your MUX, port, or interface card.

Other Terminal Problems

The other type of problem you're likely to run into with terminals is that of garbage on the screen. Garbage on the screen comes in two types: garbage intermixed with valid data characters and complete garbage.

What to Check for When Garbage is Mixed with Valid Data

The following is a list of possible reasons for garbage characters intermixed with your valid data:

- Noise on the data line:
 - RS-232 Cable too long (maximum recommended length is 50 feet)
 - Data cable near electrically noisy equipment (motors, etc.)
 - Partially shorted or broken wires within the cable
 - Noisy connection (if using phone lines)
- Hardware problem with a modem, interface card, or the terminal itself
- The program performing I/O could be sending the garbage
- The Display Functns* feature of your terminal is enabled (which displays characters that would not normally print)

What to Check for When Everything Printed is Garbage

One of the most common reasons for total garbage on the screen (and certainly the *first* thing you should check) is a baud-rate mismatch. If your terminal's speed setting is different from that of the line (as set with the `stty` command), you will get garbage on your screen (if anything at all).

Here is a list of other possible reasons for total garbage on your screen.

If you have not yet logged in, try pressing the **Break** key. This tells `getty` to try the next entry in the `/etc/gettydefs` file. The `gettydefs` file can be set up so that, as `getty` tries various entries, it will also be trying various speed settings (this is usually how it's set up). `getty` will then try various speeds (with each press of the **Break** key). When the correct speed is matched, you will get a login prompt that is readable.

- The shell environment variable called `TERM` isn't set to a value appropriate to your terminal. If you have an HP terminal, try setting the value of `TERM` to `hp` (lowercase) using your shell's `set` command.
- A running process is producing garbage output
- A miswired cable
- Excessive noise on the data line
- A hardware failure (bad interface card, modem, MUX, etc.)

10 Configuring the Kernel

With each successive release of HP-UX, system administrators have increasing ability to make changes to the configuration of the HP-UX kernel without experiencing costly and inconvenient downtime. Innovations such as Dynamic Kernel Tunables and Dynamically Loadable Kernel Modules allow critical maintenance tasks to be performed without sacrificing application availability.

With these innovations comes the need for a simpler and more comprehensive mechanism to manage kernel configurations. This chapter describes a suite of kernel configuration management commands and a web- and text-based interface that provide unified kernel configuration management. It is intended for use by HP-UX system administrators.

Kernel Configuration Features

The suite of kernel configuration tools provides several key features for system administrators:

- All kernel configuration tasks can be performed in a single graphical interface.
- All kernel configuration tasks can also be performed with a cohesive set of commands with the same user interface and same behavior.
- Kernel configurations can be saved and restored, and moved between systems.
- Administrators can save any number of kernel configurations, and can switch between them at will — often without a reboot.
- The running kernel configuration is automatically backed up (if desired) before each configuration change.
- The system automatically maintains a detailed log file of all kernel configuration changes.
- Kernel modules and kernel tunable parameters now have descriptions associated with them. Kernel tunable parameters have online documentation, and descriptions of the relationships between them.
- All kernel configuration commands can produce output in both user-friendly and script-friendly formats. HP supports release-to-release compatibility for the script-friendly formats.

What Is a Kernel Configuration?

Logically, a kernel configuration is a collection of all of the administrator choices and settings needed to determine the behavior and capabilities of the HP-UX kernel. The collection includes:

- A set of kernel tunable parameter value assignments
- A set of kernel modules, each with a desired state
- A name and optional description of the kernel configuration

Physically, a kernel configuration is a directory under /stand that contains the files needed to realize the specified behavior. The directory includes:

- An HP-UX kernel executable
- A set of HP-UX kernel module files
- A kernel registry database, containing all of the above settings
- A system file, describing the above settings in human-readable form
- Various other implementation-specific files

In addition to the configuration of the running kernel, HP-UX systems can have any number of saved kernel configurations, limited only by the disk space available in /stand.

Overview of Kernel Configuration Commands

There are three primary commands used to manage kernel configurations: `kconfig`, `kmodule`, and `kctune`; and two other commands: `kcpath` and `kclog`.

The `kconfig` command manages whole kernel configurations. It allows configurations to be saved, loaded, copied, renamed, deleted, exported, imported, etc. It can also list existing saved configurations and give details about them. For more information, see “[Managing Saved Configurations with kconfig](#)” (page 189) or the `kconfig(1M)` manpage.

The `kmodule` command manages kernel modules. Kernel modules can be device drivers, kernel subsystems, or other bodies of kernel code. Each module can be unused, statically bound into the main kernel executable, or dynamically loaded. The `kmodule` command will display or change the state of any module in the currently running configuration or any saved configuration. For more information, see “[Managing Kernel Modules with kmodule](#)” (page 160) or the `kmodule(1M)` manpage.

The `kctune` command manages kernel tunable parameters. These are variables that control the behavior of the kernel. They have many uses; common ones include controlling the allocation of system resources and tuning aspects of kernel performance. The `kctune` command will display or change the value of any tunable parameter in the currently running configuration or any saved configuration. For more information, see “[Managing Kernel Tunable Parameters with kctune](#)” (page 170) or the `kctune(1M)` manpage.

The `kcp` command prints information about the location of the currently running kernel. It is intended for use by scripts and applications that need this information. See the *kcp(1M)* manpage for details.

The `kclog` command searches the kernel configuration log file. For details, see “[The Kernel Configuration Log File](#)” (page 195) or the *kclog(1M)* manpage.

Finally, users of the `mk_kernel` command, present in previous HP-UX releases, should be aware that the command can still be used. It is included as a small shell script that invokes the `kconfig` command. This older command is obsolescent and will be removed in a future release. See *mk_kernel(1M)*.

Overview of HP SMH for Kernel Configuration

You can configure and manage the kernel without remembering the syntax of the kernel configuration commands or the exact names of modules and tunables by using HP SMH, the web- and text-based HP-UX kernel configuration tool to configure and manage the kernel of your system. HP SMH has the following features:

- Web-based and text-based interfaces.
- Kernel tunable management: monitor and modify.
- Alarm management: add, modify and remove.
- Kernel module state management: modify.
- Access to manpages for tunables.
- Command preview – When a tunable, module or alarm is modified, you can use the command preview feature by choosing the **Preview** button. This will show the kernel configuration command invocation that will perform the requested task.

You can access Kernel Configuration in any of the following ways:

- From the command line with the `kcweb -t` command.
- With a web browser through the Kernel Configuration area of HP-UX System Management Homepage. See [Figure 10-1 \(page 156\)](#).

By default, the `kcweb` command invokes the Mozilla web browser. If you want to invoke `kcweb` with any other browser, set the `BROWSER` environment variable to the path name of the browser you wish to use. For more details, see the *kcweb(1M)* manpage.

Figure 10-1 Sample Kernel Configuration Web Page

The screenshot shows a Microsoft Internet Explorer window titled "Device Home Page - Microsoft Internet Explorer provided by Hewlett-Packard". The address bar shows "https://hprdc185.example.com:2381/". The main content area is titled "System Management Homepage (hprdc185)". A banner at the top right says "User: suallanp Home | Sign Out". Below the banner, a navigation menu includes "Home", "Settings", "Tasks", "Tools", "Logs", and "Help". The current page is "Home > System Configuration > Kernel Configuration". The main content is titled "Kernel Configuration" and features a table titled "1 - 50 of 182 Tunables". The table has columns: Tunable, Tuning Capability, Current Value, Next Boot Value, Default Value, Current Usage, and Module. The "Tunable" column lists entries like NSTREVENT, NSTRPUSH, NSTRSCHED, STRCTL SZ, STRMSGSZ, acctresume, and acctsuspend. The "Module" column shows categories like hpstreams and pm_acct. A search bar labeled "Search:" is at the top of the table. To the right of the table, there is a vertical scroll bar and a context menu with options such as "View Pending Changes and Reboot", "View Change Log", "Stop Usage Monitoring", "Help Overview", and "Tunables Tab Help". At the bottom of the table, there are links for "Clear selection" and "Printable View of Table". A "Display:" dropdown set to "50 Rows" is also present. A note below the table says "Select a tunable from the list pane to view its detailed information."

Other Kernel Configuration Operations

Other sections below describe some special kernel configuration operations and special uses of the kernel configuration commands.

The usage of some kernel resources can be monitored, with alarms delivered when usage rises above a set threshold. These alarms can be configured and reviewed using the `kcalarm` command or the HP SMH tool. The resource usages can be reviewed using the `kcuusage` command or HP SMH. For more information, see “Monitoring Kernel Resource Usage” (page 180).

Administrators of older versions of HP-UX may be accustomed to using text files (“system files” or “dfiles”) to specify kernel configurations and make changes to them. The format of these files has been enhanced to accommodate new kernel configuration innovations, while retaining the usefulness of a text file for configuration operations. They are particularly useful when using the same configuration on multiple systems, since they can be easily moved between systems. The use of system files is described in “Managing Configurations with System Files” (page 192). The system file formats from previous releases of HP-UX are still accepted.

All kernel configuration changes made using the kernel configuration commands are logged to the file `/var/adm/kc.log`. Details about this log file can be found in “The Kernel Configuration Log File” (page 195), and the `kconfig(5)` and `kclog(1M)` manpages.

The primary kernel configuration commands support a specialized output format that is designed for use by scripts and applications that need to parse the output of the commands. Such scripts and applications must use this specialized output format since HP does not guarantee release-to-release compatibility for any other output format of these commands. More detail is available in “Parsing Command Output” (page 196) and the `kconfig(5)` manpage.

It is possible to have an undesirable, or even unbootable, kernel configuration because of mistaken configuration changes, hardware failures, or software defects. Mechanisms exist both to prevent such problems and to help recover from them. For more details see “Recovering from Errors” (page 197).

Common Behavior for Kernel Configuration Commands

Because the kernel configuration commands are part of a unified suite, they share behavior whenever possible. Shared behaviors include command line options, output formats, exit status codes, security constraints, and persistence of changes.

Common Command Line Options

Table 10-1 lists the options shared by the kernel configuration commands `kconfig`, `kcmodule`, `kctune`, and `kclog`.

Table 10-1 Common Kernel Configuration Command Line Options

Option	Description	k c o n f i g	c m o d u l e	k c t u n e	k c l o g
-a	(all) Include all information in the output that is normally omitted for brevity.	o	o	o	
-b	(backup) Specify whether or not to update the automatic backup configuration before the requested change and specify the default backup behavior for future changes.	o	o	o	
-c	(configuration) Specify the saved configuration to manage. If omitted, manage the currently running configuration.		o	o	o
-C	(comment) Include a comment in the kernel configuration log file entry associated with this command invocation.	o	o	o	o
-d	(description) Display descriptions of each item.	o	o	o	

Table 10-1 Common Kernel Configuration Command Line Options (continued)

Option	Description	k c o n f i g	c m o d u l e	k c t u n e	k c l o g
-D	(difference) Display only elements for which there is a change being held for next boot.	o	o	o	
-h	(hold) Hold the requested changes for next boot.	o	o	o	
-P	(parse) Use the special parsable output format.	o	o	o	
-S	(set) Display only elements that have been set to something other than default.	o	o	o	
-v	(verbose) Display items using verbose output format.	o	o	o	

Common Output Formats

When retrieving information, the primary kernel configuration commands produce output in three basic output formats: table, verbose, and parsable.

By default, the commands produce a short table format. This is a format that gives one line for each item being described. Only the most commonly used information is included, in order to allow the output to fit on one line on most terminals.

With the -v (verbose) option, the commands produce a verbose output format. This format gives all available information for each item being described, taking multiple lines to do so. A blank line separates the items in the output.

With the -P (parse) option, the commands produce an output format designed to be parsed by scripts or applications. This format is described in “[Parsing Command Output](#)” (page 196). Scripts and applications must parse this output format, because HP supports release-to-release compatibility of output format only when the -P option is used.

The kernel configuration commands all use a common format for error, warning, note, and progress messages. It is the same format used by the Software Distributor package, and therefore already familiar to most administrators.

ERROR: This is an error message. It explains why the requested operation cannot complete.

WARNING: This is a warning message. The requested operation completed, but not smoothly. A situation may exist that needs correction.

- NOTE : This is a note. It provides information about how the operation completed, or other information of potential interest to the administrator.
- * This is a progress message. It displays the steps completed during the operation.

Common Exit Status Codes

All of the kernel configuration commands exit with one of the following status codes.

- 0 The operation was successful.
- 1 The requested changes could not be applied to the currently running system. They are being held and will be applied at next boot.
- 2 The operation could not complete successfully.

Common Security Constraints

Any user can run the kernel configuration commands to query configuration information. However, access to configuration information is subject to standard UNIX file system permissions on the relevant files.

Superuser privileges are required to make any configuration changes.

Persistence of Changes

By default, the kernel configuration tools will apply configuration changes to the currently running system, causing an immediate change in behavior. System administrators can override this default by specifying the -h (hold) option on any of the commands. This option causes the changes to be held until the system is rebooted. HP recommends that this option be used only when the next reboot is expected to happen soon. If the reboot doesn't happen for months after the change, the change could come as an unwelcome surprise to an administrator who has forgotten the request.

Some configuration changes cannot be applied without a reboot. These changes will be held until the system is rebooted even if the -h option is not specified. In these cases, a warning message is printed.

If multiple configuration changes are requested in a single invocation of one of the kernel configuration commands, and any one of those changes requires a reboot, all of the requested changes will be held until the system is rebooted. In particular, if a saved kernel configuration is loaded using `kconfig -l` (load), and that configuration cannot be used without a reboot, the state of the running system is not changed and the specified kernel configuration will be used at next boot instead.

Changes being held for next boot can be listed using the -D (differences) option on the `kmodule`, `kctune`, or `kconfig` commands.

Changes being held for next boot are discarded as follows: when the currently running configuration is replaced using `kconfig -i` (import), `kconfig -l` (load), or `kconfig`

`-n` (next boot); when explicitly discarded using `kconfig -H` (unhold); or when subsequent changes are made that override them. For example, if you run.

```
# kctune -h nproc=5000 # set to 5000, hold for next boot  
# kctune nproc=6000 # set to 6000, now
```

the value of `nproc` at next boot will be 6000. The change to 5000 is discarded. A warning will be printed in these situations.

Changes that are made to the currently running system are retained when the system is rebooted. They remain in effect until changed.

Managing Kernel Modules with `kmodule`

The `kmodule` command queries and changes the states of kernel modules in the currently running configuration or in a saved configuration. The HP-UX kernel is built from a number of modules, each of which is a device driver, kernel subsystem, or some other body of kernel code. A typical kernel has 200-300 modules in it.

Getting Information About Modules

When you run `kmodule` with no options, it shows you the modules on your system, their current state, and the state they will have at next boot. On a typical system, you will see many modules in static state; some modules that are unused, which are often device drivers for hardware your system doesn't have; and a handful of modules in loaded state. The states are described below.



NOTE: The `kmodule` options are listed in [Table 10-1 \(page 157\)](#).

When you use the `-c` (configuration) option, `kmodule` displays the module information from a saved configuration instead of the currently running system.

The output of `kmodule` can be varied with several options. To control which modules are listed, use the `-a` (all), `-D` (differences), and/or `-S` (set) options. The `-a` option adds required modules to the output (normally they are omitted). The `-D` option restricts the output to only those modules whose state at next boot is different from their current state. The `-S` option restricts the output to modules whose state has been explicitly set (that is, it omits required modules, unused modules, and modules added to satisfy a dependency). The output can also be restricted by listing module names on the command line.

To control the output format, use the `-d` (description), `-v` (verbose), or `-P` (parse) options.

With No Options

With no options, the output looks like this:

```
# kmodule KeyboardMUX PCIToPCI autofs cachefs cifs rng vxportals wsio  
Module      State    Cause      Notes  
KeyboardMUX  unused
```

```

PCItoPCI      static depend
autofs        static best
cachefs       auto   best    auto-loadable, unloadable
cifs          auto   explicit auto-loadable, unloadable
rng           loaded explicit loadable, unloadable
vxportal      static explicit auto-loadable, unloadable
wsio          static required

```

With the -d Option

The -d option adds the description of each module.

```
# kmod -d PCItoPCI cachefs cifs rng
Module      State   Cause      Notes
Description
PCItoPCI    static  depend
              Generic PCI to PCI Bridge Adapter CDIO
cachefs     auto   best    auto-loadable, unloadable
              Cache File System
cifs        auto   explicit auto-loadable, unloadable
              CIFS Client Module
rng         loaded explicit loadable, unloadable
              Strong Random Number Generator
```

With the -v Option

The -v option gives verbose, multiline information about each module:

```
# kmod -v autofs cachefs
Module          autofs (1.0)
Description      Automounter File System
Timestamp       Tue Sep 12 21:53:28 2006 [45078EC8]
State           static (best state)
State at Next Boot static (best state)
Capable         static unused
Depends On     module nfswrp:0.0.0
                interface HPUX_11_31_PERF:1.0

Module          cachefs (1.0)
Description      Cache File System
Timestamp       Tue Sep 12 21:53:29 2006 [45078EC9]
State           auto (best state)
State at Next Boot auto (best state)
Capable         auto static loaded unused
Depends On     module nfswrp:0.0.0
                module dat:0.0.0
                interface HPUX_11_31_PERF:1.0
```

With the -P Option

The -P option, which is designed for use by scripts or programs, gives complete control over what information is printed. For more information, see “*Parsing Command Output*” (page 196) or the *kconfig(5)* manpage.

The special keyword ALL displays all the possible categories. Compare with the -v option.

```
# kcmodule -P ALL autofs cachefs
name      autofs
desc      Automounter File System
version   1.0
timestamp      Tue Sep 12 21:53:28 2006 [45078EC8]
state     static
cause     best
next_state      static
next_cause      best
capable    static unused
depend    module nfswrp:0.0.0
depend    interface HPUX_11_31_PERF:1.0

name      cachefs
desc      Cache File System
version   1.0
timestamp      Tue Sep 12 21:53:29 2006 [45078EC9]
state     auto
cause     best
next_state      auto
next_cause      best
capable    auto static loaded unused
depend    module nfswrp:0.0.0
depend    module dat:0.0.0
depend    interface HPUX_11_31_PERF:1.0
```

Use a comma-separated list with the -P option to display the categories you want.

```
# kcmodule -P name,desc,version autofs cachefs
name      autofs
desc      Automounter File System
version   1.0

name      cachefs
desc      Cache File System
version   1.0
```

Interpreting Module Information

Looking at the sample output in “[Getting Information About Modules](#)”, you can see that each module has a name and a textual description. Each module also has a version, which typically looks like 1.0.

A kernel configuration can only use one version of any given module. However, multiple versions may be listed if, for example, your currently running system is using a different version of a module from the one that will be used at next boot. Version numbers are normally omitted from the short listing, but will be included if there’s more than one version of a module.

Each kernel module in the currently running configuration has a state, which describes how the module is being used. The possible states are:

- unused** The module is installed on the system but not in use.
- static** The module is statically bound into the kernel executable. This is the most common state. Moving a module into or out of this state requires relinking the kernel executable and rebooting.
- loaded** The module is dynamically loaded into the kernel. Newer modules support this state. Such modules may be added to the kernel configuration or removed from it without rebooting.
- auto** The module will be dynamically loaded into the kernel when it is first needed, but it hasn't been needed yet.

When `kcmodule` is giving information about the currently running system, and there are configuration changes being held for next boot, `kcmodule` will list both the current state and the state at next boot. For next boot, the same states are used, with complementary meanings:

- unused** The module will not be used.
- static** The module will be statically bound into the kernel executable.
- loaded** The module will be dynamically loaded into the kernel during the boot process.
- auto** The module will be dynamically loaded into the kernel when it is first needed after each boot.

When `kcmodule` is giving information about a saved configuration, the same states are used.

Next to each module state is a Cause that tells why the module is (or will be) in that state. The causes are:

- explicit** The system administrator explicitly chose the state.
- best** The system administrator chose to use the module, but didn't choose a specific state, so the module is in its best state as determined by the module developer.
- auto** The module was in auto state, and was automatically loaded when something tried to use it.
- required** The module was marked required by its developer.
- depend** The module is in use because some other module in the configuration depends on it.

Different modules can support different states. Nearly all modules can be in static state, but only a few support loaded or auto states. Many modules can be in unused state, but required modules cannot. The Capable line in the output shows which states a module supports.



TIP: To see if a module is required, look to see whether `unused` appears on the `Capable` line. If it does, the module is not required.

Modules often have dependencies between them. For example, device drivers typically cannot be configured into the kernel unless the driver support modules are also configured. Dependencies like this are shown on the `Depends On` lines in the output. A module can be dependent on a particular other module, specified by name and version. A module can also be dependent on an interface that must be supplied by some other module, without saying specifically which modules supply that interface. Modules that supply such interfaces have an `Exports` line in the output, listing the interfaces they export.

Changing Module States

To change the state of a module, put module state assignments on the `kmodule` command line. (Also see “[Managing Configurations with System Files](#)” (page 192).) For example, to load the CD File System module, named `cdfs`:

```
# kmodule cdfs=loaded
```

In fact, `loaded` is the developer-chosen best state for `cdfs`, so this is the same as:

```
# kmodule cdfs=best
```

To unload it:

```
# kmodule cdfs=unused
```

See the `kmodule(1M)` manpage for details.

When you change a module state using a command as in the above examples, the change will be made immediately to the currently running system, if possible. Sometimes it's not possible to make the change immediately; for example, there might be a CD file system mounted, in which case `cdfs` can't be unloaded. In those cases, `kmodule` will hold the change and apply it at next boot. A change that moves a module into or out of `static` state can never be applied immediately, and will always be held for next boot. If any change on the `kmodule` command line has to be held for next boot, they all will be.



CAUTION: Unloading an I/O interface device driver assigned to a critical resource, such as a boot disk driver, is not immediate but is held for next boot. Removing a boot disk driver will result in an unbootable system. You can run `kmodule -D` to see what modules will be removed at next boot. Run `kconfig -H` to discard all changes pending for the next boot.

When modules are moved into or out of `static` state, the `kmodule` command will run for quite a while. This is because such changes require that the kernel executable be relinked. If you have multiple such changes to make, it is best that you list them all

on the same `kcmodule` command line, or make the changes in a system file and import it. (See “[Managing Configurations with System Files](#)” (page 192).) Either of these techniques will ensure that the kernel executable is only relinked once.

Sometimes you may want to force a change to be held for next boot, instead of applying it immediately. In these cases you can use the `-h` (hold) option with `kcmodule` to force that behavior. HP recommends that this option be used only when the next boot is expected to be soon. If, for example, the next boot doesn’t happen for months after making such a change, the system administrator could be unpleasantly surprised at the effect of a pending change that had been forgotten.

Changes to saved kernel configurations can be made by using the `-c` (configuration) option. Such changes are made to the saved configuration immediately, but they won’t affect the running system until that saved configuration is either loaded or booted. See “[Managing Saved Configurations with kconfig](#)” (page 189) for more information.

When changing module states, `kcmodule` supports the `-b` (backup) option to specify backup behavior, and the `-C` option to specify a log file comment. See “[Recovering from Errors](#)” (page 197) and “[The Kernel Configuration Log File](#)” (page 195) for details.

Managing Kernel Modules with HP SMH

HP SMH can be used to query and change the states of kernel modules in the currently running configuration. Using HP SMH, you can

- Determine which modules are currently running in the kernel
- View details about a module
- Modify the state of a module

You can view the modules pane by selecting the **Modules** tab on the **HP SMH Kernel Configuration** page, as shown in Figure 10-2.

Figure 10-2 Kernel Configuration Modules Tab

The screenshot shows a Microsoft Internet Explorer window titled "Device Home Page - Microsoft Internet Explorer provided by Hewlett-Packard". The address bar shows "https://hprdc185.example.com:2381/". The main content area displays the "System Management Homepage (hprdc185)". A navigation bar at the top includes "Home", "Settings", "Tasks", "Tools", "Logs", and "Help". Below it, a breadcrumb trail shows "Home > System Configuration > Kernel Configuration". The main content is titled "Kernel Configuration" and features three tabs: "Tunables" (selected), "Modules" (highlighted with a blue border), and "Alarms". A search bar labeled "Search:" is present. A table titled "1 - 50 of 320 Modules" lists kernel modules with columns: Module, Dynamic, Modifiable, Current State, Next Boot State, and Timestamp. The table includes rows for DeviceFileSystem, DikmDrv, KeyboardMUX, LegacyDeviceDriver, MouseMUX, OOCdio, OocCore, and OneDrive. To the right of the table is a vertical toolbar with links: "View Pending Changes and Reboot", "View Change Log", "Stop Usage Monitoring", "Help Overview", and "Modules Tab Help". At the bottom of the table are buttons for "Clear selection" and "Printable View of Table". A "Display:" dropdown set to "50 Rows" is also present. A message below the table says "Select a module from the list pane to view its detailed information." The status bar at the bottom shows "Done" and "Trusted sites".

Getting Information about Modules

To get more detailed information about a particular module, do the following:

- Scroll to the module in the list.
- Click the radio button.

The module details are displayed below the list, as shown in Figure 10-3.

Figure 10-3 Kernel Configuration Module Details

The screenshot shows a Microsoft Internet Explorer window titled "Device Home Page - Microsoft Internet Explorer provided by Hewlett-Packard". The address bar shows "https://hprdc185.example.com:2381/". The main content area is titled "System Management Homepage (hprdc185)". A navigation bar at the top includes "Home", "Settings", "Tasks", "Tools", "Logs", "Help", and "User: suallanp Home | Sign Out". Below the navigation bar, the path "Home > System Configuration > Kernel Configuration" is displayed. The main content is titled "Kernel Configuration" and features three tabs: "Tunables" (selected), "Modules" (highlighted in blue), and "Alarms". A search bar is present above a table of modules. The table has columns: Name, State, Parameters, and Last Modified. The "cifs" module is selected and highlighted in blue. To the right of the table, there is a context menu with options like "View Pending Changes and Reboot", "View Change Log", "Stop Usage Monitoring", "Help Overview", and "Modules Tab Help". Below the table, there are links for "Clear selection" and "Printable View of Table". A "Display:" dropdown is set to "50 Rows". At the bottom of the page, there is a "Details for Module : cifs" section containing a table of module details. The "cifs" module details table has columns: Description, Timestamp, Current State, Current Cause, Next Boot State, Next Boot Cause, Capabilities, Dynamic, Required, and Dependencies. The "Dependencies" row lists "module nts_cclient:0.0.0, interface HPUX_11_31_PERF:1.0". The bottom of the page includes standard browser controls like Back, Forward, Stop, Refresh, and a Trusted sites link.

Description	CIFS Client Module
Timestamp	Tue Nov 14 16:07:19 2006
Current State	auto
Current Cause	explicit
Next Boot State	auto
Next Boot Cause	explicit
Capabilities	auto static loaded unused
Dynamic	yes
Required	no
Dependencies	module nts_cclient:0.0.0, interface HPUX_11_31_PERF:1.0

Interpreting Module Information

The module details block and the module list contain the information shown in Table 10-2.

Table 10-2 Kernel Module Details Fields

Field Name	Description
Module	The name of the module
Description	A brief description of the module
Timestamp	The version of the module
Current State	The state of the module in the kernel that is currently running (unused, static, loaded, auto)

Table 10-2 Kernel Module Details Fields (continued)

Field Name	Description
Current Cause	The reason why the module is in its current state (<code>explicit</code> , <code>auto</code> , <code>depend</code> , <code>required</code> , <code>default</code>)
Next Boot State	The state of the module after the system is restarted.
Next Boot Cause	The reason why the module is in its next boot state.
Capabilities	All the states that the module is capable of supporting.
Dynamic	Is it a dynamically loadable kernel module?
Required	Does the kernel require the module?
Dependencies	All the modules required by this module.
Exports	All the interfaces exported by this module.
Modifiable	Can the module be modified?

Changing Module States

To change the state of a module, do the following:

1. Select the module.
2. Click **Modify Module** in the right-hand column of the window.




NOTE: If the cause is dependent or required, **Modify Module** will not appear, as HP SMH does not allow modifications to the state of a required module or a module on which other modules are dependent.

The **Modify Kernel Module** window is displayed, as shown in Figure 10-4.

Modify Kernel Module

Figure 10-4 Modify Kernel Module



The **Modify Kernel Module** display fields are described in Table 10-3.

Table 10-3 Modify Kernel Module Fields

Field Name	Description
Module Name	Name of the module to be modified
Description	Description of the module
Version	Version number of the module
Current State	Current value of the module
Current Cause	How the module got into its current state
Capabilities	All the states that the module can support

Table 10-3 Modify Kernel Module Fields (continued)

Field Name	Description
Dynamic	Whether the module is a dynamically loadable kernel module
Dependencies	All the modules on which this module depends

3. This window also displays the areas that you can change, as shown in Table 10-4. The areas displayed depend on the capabilities of the module. Enter your changes.

Table 10-4 Modify Kernel Module Change Options

Field Name	Description
Next Boot State	Choose the new state for the module. Select one of <code>auto</code> , <code>static</code> , <code>loaded</code> , or <code>unused</code> . The default is <code>auto</code> .
Mode of Change	Select either change immediately or change at next boot . The default is change at next boot .
Backup	Check whether to back up the current configuration before applying the change. By default, this check box is selected.
Reason for Change	Enter comments about the change in the module state

4. (Optional) To see the command that will execute the changes, click the **Preview** button.
5. After you have entered your changes, click the **Modify** button to execute them. Or click the **Cancel** button to discard your changes.

Managing Kernel Tunable Parameters with kctune

The `kctune` command is used to query and change the values of kernel tunable parameters (“tunables”), in the currently running configuration or in a saved configuration. Tunables are variables that govern the behavior of the HP-UX kernel. Tunables are used for a variety of different tasks: some control resource allocations; others control security policies; others enable optional kernel behavior; etc. There are 150-200 tunables in a typical kernel. See the `kctune(1M)` manpage.

System administrators can create their own “user-defined” tunables if they choose. These will not affect the operation of the system directly, but they can be used in computing the values of other tunables. For example, an administrator could choose to create a `num_databases` tunable, and then set several kernel tunables based on its value. A subsequent change to the value of `num_databases` would cause all of the related kernel tunable values to be changed as well.

Getting Information About Tunables

When you run kctune with no options, it shows you the tunables associated with the kernel modules on your system (as well as any user-defined tunables), their current values, and the expressions used to compute those values. If there are changes to those values being held for next boot, those will be shown as well. On a typical system, the expression for most tunables is Default, meaning that the administrator is allowing the system to choose the tunable value.

When you use the -c (configuration) option, kctune displays the tunable information from a saved configuration instead of the currently running system.

The output of kctune can be varied with several options. To control which tunables are listed, use the -D (differences) or -S (set) option. The -D option restricts the output to only those tunables whose value at next boot is different from their current value. The -S option restricts the output to only those tunables that are set to a nondefault value. The output can also be restricted by listing tunable names on the command line.

To control the output format, use the -d (description), -g (group), -v (verbose), or -P (parse) option.

With No Options With no options, the output looks like this:

```
# kctune acctresume maxuprc nproc
Tunable      Value   Expression   Changes
acctresume    4       Default
maxuprc      256     Default     Immed
nproc        4200    Default     Immed
```

With the -d Option The -d option adds the description of each tunable:

```
# kctune -d acctresume maxuprc nproc
Tunable      Value   Expression   Changes
Description
acctresume    4       Default
          Relative percentage of free disk space required to resume accounting
maxuprc      256     Default     Immed
          Maximum number of processes for each non-root user
nproc        4200    Default     Immed
          Maximum number of processes on the system
```

With the -g Option The -g option adds the name of the module defining the tunable, and sorts the output by module name. This has the effect of grouping related tunables together in the output.

```
# kctune -g acctresume maxuprc nproc
Module    Tunable      Value   Expression   Changes
pm_acct  acctresume    4       Default
pm_proc  maxuprc      256     Default     Immed
pm_proc  nproc        4200    Default     Immed
```

With the -v Option The -v option gives verbose, multiline information about each tunable:

```
# kctune -v acctresume maxuprc nproc
Tunable      acctresume
Description   Relative percentage of free disk space required to resume accounting
```

```

Module          pm_acct
Current Value  4 [Default]
Value at Next Boot 4 [Default]
Value at Last Boot 4
Default Value   4
Constraints      acctresume >= -100
                  acctresume <= 101
                  acctresume > acctsuspend
Can Change     At Next Boot Only

Tunable         maxuprc
Description    Maximum number of processes for each non-root user
Module        pm_proc
Current Value 256 [Default]
Value at Next Boot 256 [Default]
Value at Last Boot 256
Default Value  256
Constraints      maxuprc >= 3
                  maxuprc <= nproc - 5
Can Change     Immediately or at Next Boot

Tunable         nproc
Description    Maximum number of processes on the system
Module        pm_proc
Current Value 4200 [Default]
Value at Next Boot 4200 [Default]
Value at Last Boot 4200
Default Value  4200
Constraints      nproc >= 100
                  nproc <= 131072
                  nproc >= semmnu + 4
                  nproc >= maxuprc + 5
                  nproc <= nkthread - 100
Can Change     Immediately or at Next Boot

```

With the -P Option The **-P** option, which is designed for use by scripts or programs, gives you complete control over what information is printed. For more information, see “*Parsing Command Output*” (page 196) or the *kconfig(5)* manpage.

The special keyword **ALL** displays all the possible categories. Compare with the **-v** option.

```
# kctune -P ALL nproc
name      nproc
module    pm_proc
desc      Maximum number of processes on the system
defvalue  4200
bootvalue 4200
current  4200
next_boot 4200
expr     Default
next_expr Default
min      100
max      131072
dynamic  y
canauto n
default  y
auto_default  n
next_default  y
signed   n
flags    0x6c3
constraint  nproc >= 100
constraint  nproc <= 131072
```

```
constraint      nproc >= semmnu + 4
constraint      nproc >= maxuprc + 5
constraint      nproc <= nkthread - 100
```

Use a comma-separated list with the `-P` option to display the categories you want.

```
# kctune -P name,current maxuprc nproc
name    maxuprc
current 256

name    nproc
current 4200
```

Interpreting Tunable Information

Looking at the sample output above, you can see that each tunable has a name and a textual description. Each tunable is associated with a kernel module whose name is listed in the verbose output (or in the table output if `-g` is specified). Tunables can be seen and changed only if they are associated with a module that is installed on the system (or are user-defined). The module does not have to be in use.

When displaying tunable information for the currently running system, `kctune` includes the current tunable value and the expression used to compute it. If changes to the tunable's value are being held for next boot, the next boot value and expression are also shown. Verbose listings also show the value the tunable had when the system was last booted. When displaying tunable information for a saved configuration, `kctune` displays only a current value.

Tunable values are computed integer expressions, which can refer to other tunable values. (Circular references are not permitted.) The value of a tunable could be 4200, or 0x400, or 12*1024, or 4*nproc+20. Values and expressions use the syntax of the C programming language. Therefore, numbers can be written in decimal (256), octal (01000), or hexadecimal (0x100). Expressions can use the following operators and symbols:

```
( ) ~ ! - + * / % << >> < <= > >= & ^ | == != && || ?:
```

A few tunables also support values specified as percentages, for example, 10%.

White space is not permitted in any tunable expression. For backward compatibility, tunable names used in expressions can appear in all capitals, but this usage is discouraged and support for it will be removed in a future release.

All kernel tunables have a default value, which is chosen by the developer, and is shown in the verbose output. For some tunables, the default value is fixed and never changes. For other tunables, a new default value is chosen by the system at boot time. Still others can be automatically tuned, which means that the default value can change periodically while the system is running, in response to changing system resources and needs. When a tunable is set to default, its expression is reported as Default, as seen in the examples above. In these cases, the system is free to choose the value it

thinks optimal, and to change it as needed. HP recommends that tunables be left set to default unless the default is known to be unsatisfactory.



NOTE: Setting a tunable to Default is not the same thing as setting it explicitly to the default value reported by kctune. Using the example above, if you set nproc to 4200, its value will remain 4200 until you change it. However, if you set nproc to Default, its value will be kept up to date with any changes HP makes to the default value for nproc.

Some tunables have constraints on their values, which are shown in the verbose output. Sometimes these are minimum and/or maximum values, as shown for nproc above. Other times these are fixed relationships between tunables (for example, acctresume must be greater than acctsuspend) or restrictions on the allowed values (for example, dn1c_hash_locks must be a power of two). These constraints are enforced whenever tunable values are changed. There are other constraints, not shown by kctune, that are based on the current state of the system and can change over time (for example, nproc cannot be set to less than the number of processes currently running). These constraints are enforced only when changing the currently running system, and not when making changes held for use at next boot or changes to a saved configuration.

Some tunables have restrictions on when their values can be changed. These restrictions are noted in the kctune output. Tunables whose values can be changed immediately are marked Immed. Tunables whose values can be automatically tuned by the system are marked Auto. Tunables without either marking can only be changed with a reboot.

All HP-UX tunables have manpages. To obtain information about the behavior, allowed values, and side effects of a tunable, consult the manpage for that tunable, which can be found in section 5 of the online document. An overview of the kernel tunables can be found in *Tunable Kernel Parameters*, available on <http://docs.hp.com>.

Changing Tunable Values

To change the value of a tunable, put tunable value assignments on the kctune command line. (Or see “Managing Configurations with System Files” (page 192).) For example, to set nproc to 4300:

```
# kctune nproc=4300
```

To set a tunable to Default, either of these assignments will work. (Setting a user-defined tunable to Default causes it to be removed.)

```
# kctune nproc=
# kctune nproc=Default
```

Assignments can use expressions, as noted above. Note that the assignment may need to be quoted to avoid interpretation by the shell.

```
# kctune 'nkthread=nproc*2+100'
```

To create a user-defined tunable, use the `-u` (user-defined) option when you assign the tunable a value. The `-u` option is not needed to change the value of an existing user-defined tunable.

Using the `+=` symbol, you can increase the value of a tunable (by 100, in this example):

```
# kctune nproc+=100
```

Using the `>=` symbol, you can ensure a minimum value of a tunable. The command:

```
# kctune 'nproc>=5000'
```

will set `nproc` to 5000 if its current value is below 5000. If its current value is already 5000 or greater, it will be left unchanged. Note that the assignment is quoted to avoid interpretation by the shell.

See the `kctune(1M)` manpage for details.

When you change a tunable value using a command as in the above examples, the change will be made immediately to the currently running system, if possible. Sometimes it's not possible to make the change immediately; for example, you might be trying to reduce the maximum value of some resource to below the current usage. Also, there are some tunables that cannot be changed without a reboot. In those cases, `kctune` will hold the change and apply it at next boot. If any change on the `kctune` command line has to be held for next boot, they all will be.

Sometimes you may want to force a change to be held for next boot, instead of applying it immediately. In these cases you can use the `-h` (hold) option of `kctune` to force that behavior. HP recommends that this option be used only when the next boot is expected to be soon. If, for example, the next boot doesn't happen for months after making such a change, the system administrator could be unpleasantly surprised at the effect of a pending change that had been forgotten.

Changes to saved kernel configurations can be made by using the `-c` (configuration) option. Such changes are made to the saved configuration immediately, but they won't affect the running system until that saved configuration is either loaded or booted. See “[Managing Saved Configurations with kconfig](#)” (page 189) for more information.

When changing tunable values, `kctune` supports the `-b` (backup) option to specify backup behavior, and the `-C` option to specify a log file comment. See “[Recovering from Errors](#)” (page 197) and “[The Kernel Configuration Log File](#)” (page 195) for details.

Managing Kernel Tunable Parameters with HP SMH

HP SMH can be used to query and change the values of kernel tunable parameters (“tunables”) in the currently running configuration. Using HP SMH, you can:

- Modify the value of a tunable
- View details about a tunable
- Search for a tunable
- Check the current and next boot value for a tunable
- Print details about a tunable or print a list of all tunables

You can view the tunables pane by selecting the **Tunables** tab on the HP SMH **Kernel Configuration** page, as shown in Figure 10-5.

Figure 10-5 Kernel Configuration Tunables Tab

The screenshot shows a Microsoft Internet Explorer window titled "Device Home Page - Microsoft Internet Explorer provided by Hewlett-Packard". The address bar shows the URL "https://hprdc185.example.com:2381/". The main content area is titled "System Management Homepage (hprdc185)". A banner at the top right shows a user profile for "suallanp" with options to "Home" or "Sign Out". Below the banner, a navigation bar includes "Home", "Settings", "Tasks", "Tools", "Logs", and "Help". The current page is "Home > System Configuration > Kernel Configuration". The main content is titled "Kernel Configuration" and features a table titled "Tunables". The table has columns: Tunable, Tuning Capability, Current Value, Next Boot Value, Default Value, Current Usage, and Module. The table lists several tunables, such as NSTREVENT, NSTRPUSH, NSTRSCHED, STRCTL SZ, STRMSGSZ, acctresume, and acctsuspend. A search bar is present above the table. To the right of the table, there is a vertical toolbar with links like "View Pending Changes and Reboot", "View Change Log", "Stop Usage Monitoring", "Help Overview", and "Tunables Tab Help". At the bottom of the table, there are buttons for "Clear selection" and "Printable View of Table". A "Display" dropdown menu is set to "50 Rows". A note below the table says "Select a tunable from the list pane to view its detailed information." The bottom of the window shows standard browser controls and a "Trusted sites" icon.

Getting Information About Tunables

To get more detailed information about a particular tunable, do the following:

1. Scroll to the tunable in the list.
2. Click the radio button.

The tunable details are displayed under the list, as shown in Figure 10-6.

Figure 10-6 Kernel Configuration Tunable Details

The screenshot shows a Microsoft Internet Explorer window titled "Device Home Page - Microsoft Internet Explorer provided by Hewlett-Packard". The address bar shows the URL <https://hprdc185.example.com:2381/>. The page title is "System Management Homepage (hprdc185)". The navigation menu includes Home, Settings, Tasks, Tools, Logs, Help, and a link to "Kernel Configuration". The main content area is titled "Kernel Configuration" and contains three tabs: Tunables (selected), Modules, and Alarms. A search bar is present above the tunable list. The list displays 1 - 50 of 182 Tunables. One tunable, "fs_symlinks", is selected and highlighted in blue. To the right of the list is a context menu with options like "View Pending Changes and Reboot", "View Change Log", "Stop Usage Monitoring", "Help Overview", "Tunables Tab Help", "Modify Tunable", and "Printable View of Table". Below the list are links for "Clear selection" and "Printable View of Table". A "Display:" dropdown is set to "50 Rows". The bottom of the window shows standard browser controls and a "Trusted sites" icon.

Description	Module	Tuning Capability	Default	Current Value	Next Boot Value	Next Boot Value (Expression)	Range	Constraints
Maximum number of symbolic links used to resolve a path name	fs	Dynamic	20	<Default>	<Default>	Default	20..1024	fs_symlinks >= 20 fs_symlinks <= 1024

Interpreting Tunable Information

The tunable details block and the tunable list contain the information shown in Table 10-5.

Table 10-5 Tunable Details Fields

Field Name	Description
Tunable	The name of the tunable.
Description	A brief description of the tunable.
Module	The name of the module (if any) that the tunable is associated with.
Tuning Capability	Whether the tunable is Static, Dynamic, or Auto.
Default	The default value for the tunable.

Table 10-5 Tunable Details Fields (continued)

Field Name	Description
Current Value	The current maximum value for the resource.
Next Boot Value	The planned value, with all formulae computed.
Next Boot Value (Expression)	The formula (or integer) describing the next boot value.
Range	The range of values that are valid for the tunable.
Constraints	The dependencies that a tunable might have on other tunables as well as value restrictions.
Current Usage	The percentage of the resource that is being used.


Changing Tunable Values

To change the value of a tunable, execute the following steps:

1. Select the tunable.
2. Click **Modify Tunable** in the right-hand column of the window.

The **Modify Kernel Tunable** page is displayed, as shown in Figure 10-7.

Figure 10-7 Modify Kernel Tunable



The **Modify Kernel Tunable** page displays the fields shown in Table 10-6.

Table 10-6 Modify Kernel Tunable Fields

Field Name	Description
Tunable Name	The name of the tunable that will be modified.
Description	Description of the tunable.
Module	The kernel module that the tunable is associated with.
Default	The default value of the tunable.
Range	The range of acceptable values for the tunable. Negative numbers are indicated by a minus sign (-). Positive numbers are unsigned. Not Available indicates that the underlying command, kctune, is returning neither a minimum nor a maximum value.

Table 10-6 Modify Kernel Tunable Fields (continued)

Field Name	Description
Current Value	The current value of the tunable.
Tuning Capability	Displays whether the tunable is Static, Dynamic, or Auto.
Current Usage	The percentage of the resource that is being used.
Constraints	The dependencies that a tunable might have on other tunables and its value restrictions

3. The **Modify Kernel Tunable** page also displays the areas that you can change, as shown in Table 10-7. The areas displayed depend on the capabilities of the tunable. Enter your changes.

Table 10-7 Modify Kernel Tunable Change Options

Field Name	Description
New Setting (Expression/Value)	Enter a formula (or integer) describing the new value. You can set it to the default by clicking the Reset to Default button.
New Setting (Evaluated)	The calculated value of the new setting. You may have to click the Recalculate Value button.
Mode of Change	Select either change immediately or change at next boot . The default is change at next boot . This field only appears for dynamic tunables. Other tunables always change at next boot.
Backup	Check whether to back up the current configuration before applying the change. By default, this check box is selected.
Reason for Change	Enter comments about the change in the tunable value.

4. (Optional) To see the command that will execute the changes, click the **Preview** button.
5. After you have entered your changes, click the **Modify** button to execute them. Or click the **Cancel** button to discard your changes.

Monitoring Kernel Resource Usage

Some tunable parameters represent kernel resources whose usage can be monitored. For these tunables, you can set alarms to notify you when the usage of the corresponding kernel resource crosses a threshold you specify.

The alarms page allows you to:

- Create and remove alarms
- Activate and deactivate alarms

- Find alarms that have been triggered
- View details on alarms

You can view the alarms pane by selecting the **Alarms** tab on the HP SMH **Kernel Configuration** page, as shown in Figure 10-8.

Figure 10-8 Kernel Configuration Alarms Tab

The screenshot shows a Microsoft Internet Explorer window titled "Device Home Page - Microsoft Internet Explorer provided by Hewlett-Packard". The address bar shows "https://hprdc185.example.com:2381/". The main content area displays the "System Management Homepage (hprdc185)". A navigation bar at the top includes "Home", "Settings", "Tasks", "Tools", "Logs", and "Help". Below this is a breadcrumb trail: "Home > System Configuration > Kernel Configuration". The main content area has a title "Kernel Configuration" and a tab bar with "Tunables", "Modules", and "Alarms" (which is selected). A search bar is present. A table titled "1 - 3 of 3 Alarms" lists three tunables: "any", "filecache_max", and "nproc". Each row includes columns for Tunable, Alarm Status, Threshold, Current Usage, and Notification Type. To the right of the table is a list of links: "Add Alarm", "View Pending Changes and Reboot", "View Change Log", "Stop Usage Monitoring", "Help Overview", and "Alarms Tab Help". At the bottom of the table area, there are links for "Clear selection" and "Printable View of Table". A "Display:" dropdown is set to "50 Rows". A message "Select an alarm from the list pane to view its detailed information." is displayed below the table. The bottom of the window shows the standard Internet Explorer status bar with icons for Stop, Back, Forward, Home, and Trusted sites.

Tunable	Alarm Status	Threshold	Current Usage	Notification Type
any	on	80%	-	email
filecache_max	on	66%	10%	email
nproc	on	70%	4%	console

Getting Information about Alarms

To get more detailed information about a particular alarm, do the following:

1. Scroll to the alarm in the list.
2. Click the radio button to select it.

The alarm details are displayed under the list, as shown in Figure 10-9.

Figure 10-9 Kernel Configuration Alarm Detail

Tunable	Alarm Status	Threshold	Current Usage	Notification Type
any	on	80%	-	email
filecache_max	on	66%	10%	email
npoc	on	70%	4%	console

Details for Alarm :

Tunable	filecache_max
Description of Tunable	Maximum amount of physical memory to be used for caching file I/O data
Current Usage Value	10%
Event Type	initial, return
Alarm Status	on
Polling Interval	5 minutes
Notification Type	email
Notification Data	support@it.example.com
Comment	-
Tuning Capability	Auto

Interpreting Alarm Information

The alarm details block and the alarms list contain the information shown in Table 10-8.

Table 10-8 Kernel Configuration Alarm Fields

Field Name	Description
Tunable	The name of the tunable.
Description of Tunable	The description of the tunable.
Current Usage Value	The percentage of resource being consumed at the previous polling.
Event Type	The event notification to be used.
Alarm Status	The status of the alarm, one of the following: on The alarm is active. ringing The alarm has been triggered. off The alarm is deactivated.

Table 10-8 Kernel Configuration Alarm Fields (continued)

Field Name	Description
Polling Interval	The time interval between polling.
Threshold	The percentage at which the alarm should activate.
Notification Type	The notification method used when the alarm is triggered. The notification types are: console, email, opcmsg, snmp, syslog, textlog, tcp, and udp. See Table 10-10 for details.
Notification Data	Supplementary information used by the notification method.
Comment	The comment field; some comment data is added automatically when alarms are deactivated.
Tuning Capability	One of static, dynamic, or auto.


Changing Alarm Settings

To change the settings for an alarm, execute the following steps:

1. Select the alarm.
2. Click **Modify Alarm** in the right-hand column of the window.

The **Modify Kernel Alarm** page is displayed, as shown in Figure 10-10.

Figure 10-10 Modify Kernel Alarm



The **Modify Kernel Alarm** page displays the fields shown in Table 10-9.

Table 10-9 Modify Alarm Fields

Field Name	Description
Tunable	The name of the tunable.
Description of Tunable	The description of the tunable.
Current Usage Value	The percentage of resource being consumed at the previous polling.
Event Type	When notifications are to be sent: initial, repeat, return. See Table 10-10 for details.

Table 10-9 Modify Alarm Fields (continued)

Field Name	Description
Alarm Status	The status of the alarm, one of the following: <code>on</code> The alarm is active. <code>ringing</code> The alarm has been triggered. <code>off</code> The alarm is deactivated.
Polling Interval	The interval, in minutes, between polling of resource usage.
Notification Type	The notification method: <code>console</code> , <code>opcmsg</code> , <code>syslog</code> , <code>textlog</code> , <code>email</code> , <code>snmp</code> , <code>tcp</code> , <code>udp</code> .
Notification Data	Supplementary information used by the notification method.
Comment	The comment field; some comment data is added automatically when alarms are deactivated.
Tuning Capability	One of <code>static</code> , <code>dynamic</code> , or <code>auto</code> .

3. The **Modify Kernel Alarm** page also displays the areas that you can change, as shown in Table 10-10. The areas displayed depend on the capabilities of the alarm. Enter your changes.

Table 10-10 Modify Alarm Change Options

Field Name	Description
Threshold	Enter the percentage at which the alarm should activate.
Event Type	Check the boxes that determine when notifications are to be sent: <code>initial</code> First polling at which resource usage exceeds threshold. Also sent when the alarm is first added, activated, deactivated, or the system reboots. <code>repeat</code> Each polling at which resource usage exceeds the threshold. This can lead to a large number of messages if the polling interval is small. <code>return</code> First polling at which resource usage falls below threshold. If no box is checked, the default event type, as set by <code>kcalarm</code> , is used. Note: Selecting both <code>initial</code> and <code>return</code> will generate a notification whenever the usage crosses above or below the threshold.
Polling Interval	Enter the time interval (in minutes) between polling.

Table 10-10 Modify Alarm Change Options (continued)

Field Name	Description
Notification Type	Select the notification method to be used when the alarm is triggered. The choices are: console Send a message to the system console. email Send an e-mail to the specified address. Fill in the Email Address field. opcmsg Send messages to ITO and OpenView applications via the opcmsg daemon. Select a value for Notification Data , one of normal, warning, minor, major, critical. snmp Send messages to applications, such as Network Node Manager, that use SNMP traps. Select a value for Notification Data , one of normal, warning, minor, major, critical. syslog Log the alarm in the system log file. textlog Log the alarm in a text file. The file is stored in the directory /var/opt/resmon/log. Fill in a file name in the File Name field. tcp Send TCP encoded events to the specified target host name and port. Fill in the Host Name and Port Number fields. udp Send UDP encoded events to the specified target host name and port. Fill in the Host Name and Port Number fields.
Comment	Enter an optional comment.

4. (Optional) To see the command that will execute the changes, click the **Preview** button.
5. After you have entered your changes, click the **Modify** button to execute them. Or click the **Cancel** button to discard your changes.

Add an Alarm

To create a new alarm, execute the following steps:

1. Click **Add Alarm** in the right-hand column of the **Kernel Configuration Alarms** window.

The **Add Alarm** page is displayed, as shown in Figure 10-11.

Figure 10-11 Add Alarm

The screenshot shows a Microsoft Internet Explorer window titled "Device Home Page - Microsoft Internet Explorer provided by Hewlett-Packard". The address bar shows the URL <https://hprdc185.example.com:2381/>. The title bar says "System Management Homepage (hprdc185)". The main content area is titled "Add Alarm". It contains the following fields:

- Tunable:** A dropdown menu set to "any".
- Threshold:** A text input field containing "80 %".
- Event Type:** A group of checkboxes: "initial" (checked), "repeat" (unchecked), and "return" (unchecked).
- Polling Interval:** A text input field containing "5 minutes".
- Notification Type:** A dropdown menu set to "console".
- Comment:** An empty text area.

At the bottom of the form are four buttons: "Add", "Preview", "Cancel", and "Help". Below the form, a note says "To see the command(s) that will be run, input the required information and press the [Preview] button." The status bar at the bottom of the browser window shows "Done" and "Trusted sites".

2. In the dropdown list of the **Tunable** field, select the tunable you want to monitor. Note that only certain tunables are included in the list. If you select any, an alarm will be set on all the tunables in the list.

You can set more than one alarm on any selectable tunable.
3. Enter values for the **Threshold**, **Event Type**, **Polling Interval**, **Notification Type**, and **Comment** fields, as described in Table 10-10.
4. (Optional) To see the command that will create the alarm, click the **Preview** button.
5. After you have entered your values, click the **Add** button to create the alarm. Or click the **Cancel** button to discard it.

Activate, Deactivate, or Remove an Alarm

To activate, deactivate, or remove an alarm, execute the following steps:

1. Select the alarm.

The alarm details are displayed under the list, as shown in Figure 10-9.

- At the bottom of the window, click one of the following buttons:
 - Activate** to activate the alarm. The **Alarm Status** changes to **on**.
 - Deactivate** to deactivate the alarm. The **Alarm Status** changes to **off**.
 - Remove** to delete the alarm. The alarm is removed from the list.

Resource Usage Commands

The `kcalarm` command is used to add, delete, or list selected kernel tunable alarms, as well as turn kernel tunable monitoring on and off.

`kcalarm` is used to manage selected kernel tunable alarms and monitors; alarms and monitors are implemented in the `kcmnd` daemon. Users can create, modify, delete, and list selected kernel tunable alarms. Alarms send a notification through various notification targets when a kernel tunable crosses a specified percentage threshold of its current setting.

Usage monitoring is the process of collecting historical tunable data. When this feature is turned on, historical data is collected on the usage of supported tunables. These data are used by the `kcusage` command to generate usage tables (including top consumers) for supported kernel tunables. These data also enable usage graphs in the HP SMH tool. Monitoring is turned on by default.

For more information, see the `kcalarm(1M)`, `kcmnd(1M)`, and `kcusage(1M)` manpages.

To start or stop usage monitoring, you can click **Start Usage Monitoring** or **Stop Usage Monitoring** in the right-hand column of any tab on the **Kernel Configuration** page

Managing the Running Configuration Using `kconfig`

The `kconfig` command has two options that are useful for dealing with changes to the currently running kernel configuration that are being held for next boot.

Configuration changes are held for next boot when requested (using the `-h` (hold) option of `kmodule` or `kctune`, or the `-n` (next boot) option of `kconfig`). Configuration changes are also held for next boot when they cannot be applied to the currently running system.

To get a list of changes being held for next boot, run `kconfig -D` (differences). This is really just a short cut for running `kmodule -D` and `kctune -D`. Similarly, to get a list of configuration settings that are set to nondefault values, run `kconfig -S` (set). This is a short cut for running `kmodule -S` and `kctune -S`.

If you decide that you don't want those changes to be applied at next boot after all, run `kconfig -H` (unhold). All changes being held for next boot will be discarded.

For more information on changes being held for next boot, see “[Persistence of Changes](#)” (page 159).

Managing Saved Configurations with kconfig

When you have an HP-UX kernel configuration that satisfies your needs, you may want to save a copy of it to protect yourself against inadvertent configuration changes. Or, you may want to have multiple kernel configurations, so that you can switch between them easily. HP-UX allows you to save as many kernel configurations as you wish (subject to available disk space in /stand), and to modify them and use them at will.

Getting Information about Saved Configurations

When you run `kconfig` with no options, it shows you the saved configurations on your system. There will always be a saved configuration called `backup`, which is automatically maintained by the system; any other saved configurations on the system will also be listed. (For more information on the backup configuration, see “[Recovering from Errors](#)” (page 197).)

The output of `kconfig` can be varied with several options. The output can be restricted to specific configurations by listing them on the command line.

To control the output format, use the `-a` (all), `-v` (verbose), or `-P` (parse) options.

With No Options With no options, the output looks like this:

```
# kconfig
Configuration Title
backup      Automatic Backup
day         Configuration for daytime multiuser processing
last_install Created by last OS install
night       Configuration for nighttime batch processing
```

With the `-v` Option The `-v` option gives verbose, multiline information about each saved configuration:

```
# kconfig -v day
Configuration day
Title      Configuration for daytime multiuser processing
Created    Thu Oct 12 01:33:36 2006 by allanp
           as a copy of 'last_install'
Modified   Thu Oct 12 01:37:14 2006 by allanp
Kernel Path /stand/day/vmunix
```

With the `-a` Option The `-a` option is the equivalent of the command sequence:

```
# kcmodule -a -v -c config
# kctune -v -c config
# kconfig -v config
```

for each specified or implied configuration (`config`). This gives a detailed record of all settings in the configuration. The following output is edited.

```
# kconfig -a day
Module          DeviceFileSystem (1.0)
Description     DevFS File System
Timestamp       Mon Sep 11 15:31:18 2006 [4505E3B6]
```

```

State           unused
Capable        static unused
Depends On    module OocCore:0.0.0
                interface HPUX_11_31_PERF:1.0

...
Tunable         NSTREVENT
Description     Maximum number of concurrent Streams bufcalls
Module          hpstreams
Current Value   50 [Default]
Default Value   50

...
Configuration day
Title          Configuration for daytime multiuser processing
Created        Thu Oct 12 01:33:36 2006 by allanp
                as a copy of 'last_install'
Modified       Thu Oct 12 01:37:14 2006 by allanp
Kernel Path    /stand/day/vmunix

```

With the -P Option The -P option, which is designed for use by scripts or programs, gives complete control over what information is printed. For more information, see “Parsing Command Output” (page 196) or the *kconfig(5)* manpage.

The special keyword ALL displays all the possible categories. Compare with the -v option.

```
# kconfig -P ALL day
name      day
title    Configuration for daytime multiuser processing
createtime   Thu Oct 12 01:33:36 2006
createuser   allanp
modifytime   Thu Oct 12 01:37:14 2006
modifyuser   allanp
kernel   /stand/day/vmunix
```

Use a comma-separated list with the -P option to display the categories you want.

```
# kconfig -P name,title,modifyuser night
name      night
title    Configuration for nighttime batch processing
modifyuser allanp
```

Interpreting Saved Configuration Information

Referring to the examples above, each saved configuration has a name, title, time signatures, user signatures, and kernel location.

The name must start with a letter; contain only letters, digits, and underscores; and be at most 32 characters long. Except for the backup configuration, you choose the name for each saved configuration when you create it, and you can rename it at will.

Each saved configuration can also have a title. The title can be used to provide yourself with a longer description of the configuration's purpose or settings. It is optional.

Each saved configuration also has a pair of timestamps. The `Created/createtime` stamp indicates when the configuration was created (`kconfig -s`). The `Modify/modifystime` stamp indicates when the configuration was last changed. Each timestamp is associated with the login name of the user who performed the action.

The kernel itself is located in a subdirectory of `/stand` that has the same name as the configuration.

Associated with each saved configuration is a complete set of module state settings and tunable value settings. These can be seen by using:

```
# kmodule -c config
```

and

```
# kctune -c config
```

or by using

```
# kconfig -a config
```

Using and Modifying Saved Configurations

Creating Saved Configurations

Saved kernel configurations can be created in three ways: by saving the currently running configuration, by copying an existing saved configuration, or by reading a system file.

To save the currently running configuration, use `kconfig -s` (save). The resulting saved configuration will include any changes to the currently running configuration that are being held for next boot.

An existing saved configuration can be copied using `kconfig -c` (copy).

For information on working with system files, see “[Managing Configurations with System Files](#)” (page 192).

Using Saved Configurations

A saved configuration can be loaded using `kconfig -l` (load). This changes the configuration of the currently running kernel to match what was saved. If the configuration can be changed without a reboot, the changes will take effect immediately. Otherwise, all of the changes will be held for next boot.

Sometimes you may want to force the configuration change to be held for next boot, instead of applying it immediately. In these cases, you can mark the saved configuration for use at next boot using `kconfig -n` (next boot). HP recommends that this option be used only when the next boot is expected to be soon. If, for example, the next boot doesn't happen for months after you make such a change, the system administrator

could be unpleasantly surprised at the effect of a pending change that had been forgotten.

To find out which saved configuration is marked for use at next boot, use `kconfig -w` (which). This command also identifies the saved configuration that was most recently loaded or booted, or the system file that was most recently imported.

Modifying Saved Configurations

To modify the module state settings and tunable value settings in a saved configuration, use the `-c` (configuration) option of the `kmodule` and `kctune` commands, respectively. Saved configurations can also be changed by changing their system file and then importing it; see “[Managing Configurations with System Files](#)” (page 192).

Several options of `kconfig` allow other changes to saved configurations. The `-r` (rename) option will rename a saved configuration. (The backup configuration cannot be renamed.) The `-t` option will change the title on a saved configuration. The `-d` (delete) option will delete a saved configuration.

If a configuration has been marked for use at next boot, and you decide you want to continue using the currently running configuration instead, use `kconfig -H` (unhold) to discard all changes being held for next boot.

Managing Configurations with System Files

Every kernel configuration has a corresponding system file. A system file is a flat text file that describes all of the configuration settings in a compact, machine-readable, portable format. The format of a system file is described in detail in the *system(4)* manpage. It is an enhancement of the format used in previous releases of HP-UX; the previous formats are still accepted.

Making Configuration Changes with System Files

System files provide an alternate mechanism for kernel configuration, because configuration changes can be made by editing a system file and then telling the kernel configuration tools to apply the changes. This is the kernel configuration method most familiar to users of older versions of HP-UX.

To make configuration changes using a system file, start with the system file corresponding to the configuration you want to change.² The system automatically maintains system files for each configuration. The system file for the currently running configuration is located at `/stand/system`. The system file for any saved configuration is located at `/stand/configname/system`. If you want to create a new system file for a configuration, use the `kconfig -e` (export) command. This command takes two forms:

- Export the running configuration:
2. You will be asked to confirm your changes if the system file comes from a different configuration from the one you’re changing, or if it’s out of date with respect to the configuration you’re changing.

- ```
kconfig -e filename
```
- Export a saved configuration:  

```
kconfig -e configname filename
```
- 



**NOTE:** /stand/system, and any system file created by exporting the running configuration, always reflects any changes that are being held for next boot.

Once you have a system file, you can edit it using any text editor, making the changes you desire. After editing it, you can apply the changes with the `kconfig -i` (import) command. This command has three forms:

- Import to running configuration now:  
`kconfig -i filename`
- Import and hold for next boot:  
`kconfig -h -i filename`
- Import to saved configuration:  
`kconfig -i configname filename`

In the first form, if the changes cannot be applied to the running system, they will be held for next boot.

For backward compatibility, the `mk_kernel` command is still available to apply changes made in a system file. Note, however, that its name is no longer accurate since it will apply configuration changes without making a kernel if it can. This command has the form:

```
mk_kernel [-o target] [-s filename]
```

*filename* is the name of the system file to read; if not specified, /stand/system is used. To import to a saved configuration, *target* should be the name of the configuration. To import to the currently running system, taking effect immediately if possible, *target* should be /stand/vmunix. (Changes will be held until next boot if they cannot be applied immediately.) If *target* is omitted, the changes will be made to a saved configuration called hpxx\_test. It is not possible to import to the currently running system, forcing changes to be held for next boot, using `mk_kernel`. Use `kconfig -h -i` for that purpose.

It is important to note that the system files at /stand/system and /stand/*configname*/system are automatically recreated after every configuration change. In this process, comments in the system file are not preserved. Also, the ordering of lines in the file is not preserved. Therefore, HP recommends against putting comments in the system files. Instead, use the -C (comment) option when importing the configuration, to add your comments directly to the kernel configuration log file. (See “The Kernel Configuration Log File” (page 195).)

Most changes made in system files can be made using the kernel configuration commands, and vice versa. Here are the equivalents:

| System File Line                                           | Kernel Configuration Command                 |
|------------------------------------------------------------|----------------------------------------------|
| <code>modulename</code>                                    | <code>kcmodule modulename=best</code>        |
| <code>module modulename best</code>                        | <code>kcmodule modulename=best</code>        |
| <code>module modulename state [version]<sup>1</sup></code> | <code>kcmodule modulename=state</code>       |
| <code>(no entry for modulename)</code>                     | <code>kcmodule modulename=unused</code>      |
| <code>tunablename tunablevalue</code>                      | <code>kctune tunablename=tunablevalue</code> |
| <code>tunable tunablename tunablevalue</code>              | <code>kctune tunablename=tunablevalue</code> |
| <code>(no entry for tunablename)</code>                    | <code>kctune tunablename=default</code>      |

- 1 System files created by the kernel configuration tools always list the version number for each module. However, it is not required. Administrators adding module lines to a system file need not give version numbers.

## Uses for System Files

System files are primarily useful in the following situations.

1. They are useful for system administrators who are familiar with them from previous releases of HP-UX. If you are used to editing `/stand/system` and running `mk_kernel` to make configuration changes, it will still work.
2. System files are useful if you want to apply multiple configuration changes simultaneously. You can edit a `/stand/system` and change three tunable values and two module states, and have all of those changes take effect together when you import the system file with `kconfig -i` or `mk_kernel`. By contrast, each invocation of one of the kernel configuration commands applies changes separately (although multiple changes listed on the same configuration command line are applied together).

Applying multiple changes together is particularly valuable when modules are moved into or out of `static` state, because each command that does this will run for quite a while. This occurs because such changes require that the kernel executable be relinked. If you have multiple such changes to make, it is best that you list them all on the same `kcmodule` command line, or make the changes in a system file and import it. Either of these techniques will ensure that the kernel executable is only relinked once.

3. System files are used for copying configurations from one system to another. It is not safe to copy a kernel configuration directory from one machine to another, and HP does not support doing that. However, it is perfectly safe to export a system file from a configuration on one system, move that system file to a different system,

and import it there. This is an appropriate and effective way to ensure that two machines are running compatible configurations. (Compatible means they have the same set of kernel modules, but they may have different versions of those modules due to patch installations.)

In some cases, running compatible configurations is not enough; you need to be sure that two machines are running exactly the same configuration. In that case, use the -v (version match) option while importing the system file on the target system. This option turns on strict version checking, and the import will fail if the two machines have different versions of kernel modules installed.

## The Kernel Configuration Log File

It is often useful to know what configuration changes have been made on a system. For this purpose, the kernel configuration tools automatically maintain a log file at `/var/adm/kc.log`. This file lists every change made using the kernel configuration commands. (Some configuration changes can be made by calling kernel system calls directly. These changes are not logged.) Changes made through HP SMH, the web-based interface for kernel configuration, are logged since HP SMH uses the kernel configuration commands to make the changes.


The log file is a plain text file that you can view directly. The `kclog` command is provided for when you want to do an intelligent search of the log file, but its use is optional. (More information on the `kclog` command can be found in the `kclog(1M)` manpage.)

All of the kernel configuration commands accept a -C (comment) option when they are being used to make configuration changes. The -C option allows you to specify a comment that will be included in the log entry for your change. This can help readers of the log understand the reasons for your changes.

To add a comment to the log without making a configuration change, use `kclog -C`.

On the **Kernel Configuration** page, you can select **View Change Log** in the right-hand column to see the kernel configuration log file, as shown in Figure 10-12.

**Figure 10-12 View Change Log**



## Parsing Command Output

Improvements to HP-UX often require changes in the output formats of commands like those described here. This can be troublesome when applications or scripts have been written that parse the outputs of those commands. For this reason, each of the primary kernel configuration commands (`kcmodule`, `kctune`, and `kconfig`) has a special output format, selected with the `-P` (parse) option, designed for parsing by

applications. In addition to providing release-to-release compatibility, it is also easier to parse than human-readable output.



**CAUTION:** HP reserves the right to change the other output formats of these commands at any time. HP will not support applications and scripts that parse the output of these commands unless they use the -P option.

The -P option of each of these commands takes a list of field names, identifying the fields that the application wants to have appear in the output. The available field names are different for each command and are documented in the manpages for the commands. The list is comma-separated and cannot contain spaces. Examples are shown in the sections “Managing Kernel Modules with `kcmodule`” (page 160), “Managing Kernel Tunable Parameters with `kctune`” (page 170), and “Managing Saved Configurations with `kconfig`” (page 189).

The special field name ALL can be used to retrieve all available data. When this field name is used, the output may include fields that are not listed in the manpage. The order of fields in the output is undefined.

The output format consists of one line per field, containing the field name, a single tab character (ASCII 9), the field value, and a newline (ASCII 12). The fields are printed in the order requested for each item, with empty lines between the items.

Some fields have multiple values. In these cases, there will be one line for each value of the field, each starting with the field name in the manner described.

Some fields do not have values under some circumstances. For example, the min or max tunable fields have no meaning for tunables that have no defined limits. In these cases, no line will be printed for that field.

## Recovering from Errors

Occasionally, kernel configuration changes are made that are undesirable. Also, hardware failures and changes can ruin a previously acceptable kernel configuration. HP-UX has several mechanisms available to system administrators who need to recover from such issues, including:

- The kernel configuration log file. See “The Kernel Configuration Log File” (page 195).
- Saved configurations, including the automatically maintained backup configuration. See “The Automatic Backup Configuration” (page 198) and “Booting a Saved Configuration” (page 198).
- Fail-safe boot mode with fail-safe tunable values and module loading. See “Booting in Fail-Safe Mode” (page 199).
- Boot-time overrides of kernel tunable values.

## The Automatic Backup Configuration

The system automatically maintains a saved configuration called `backup`. Generally, any time you use the kernel configuration tools to make a change to the currently running configuration, the previous (pre-change) configuration is saved to `backup`. Therefore the backup configuration is somewhat like the “undo” command in a word processor. In these cases, if you load the backup configuration using `kconfig -l backup`, it will reverse the last change you made to the currently running configuration using the kernel configuration commands.

Some changes can be made to the currently running configuration by calling kernel system calls directly. The `backup` configuration is not updated when those changes are made.

There are cases in which you may not want this automatic backup behavior. For example, if you have made an undesirable change and are trying to fix it, you do not want the kernel configuration commands to replace a good backup configuration with the one containing your undesirable change.

The `-b off` option (keep the existing `backup`) can be given in any kernel configuration command to disable the automatic update of the `backup` configuration. When making changes using HP SMH, you can clear the `back up the current configuration before applying change` check box to disable the automatic backup behavior.

When your system first boots, the `backup` configuration mirrors the configuration that was in use before the reboot. You may not want this replaced by the first kernel configuration change you make, especially since the first kernel configuration change could be made by a startup script before you even get a login prompt.

For this reason, the first configuration changes after a boot are handled specially. Instead of automatically replacing the `backup` configuration, the kernel configuration commands will ask you whether or not to do so.<sup>3</sup> They will continue to ask, each time you make a change, until the first time you say `Yes`. From that point on, until next boot, they will automatically replace the `backup` configuration with each change as described above.

If you want to disable the automatic replacement of the `backup` configuration for a particular change, specify `-b off`. If you want to force an automatic replacement of the `backup` configuration, specify `-b on (backup)`. These options work with any kernel configuration command that makes configuration changes.

## Booting a Saved Configuration

In extreme circumstances, a mistaken configuration change can result in a kernel configuration that won’t boot. In these cases, you have two options: boot a different configuration, such as the automatic backup configuration, and/or boot in fail-safe mode (described in “[Booting in Fail-Safe Mode](#)” (page 199)).

3. If the command is being run noninteractively, such as from a startup script, the answer is assumed to be `No` for `kmodule`, `kctune`, and `kcdevice`, and `Yes` for `kconfig`.

## Integrity

To boot a saved configuration on an Integrity system, interrupt the automatic boot process when it reaches the point that it has started the HP-UX boot loader. (On most systems, this is during the second 10-second countdown.) At the **HPUX>** prompt, type

```
HPUX>boot thursday
```

## HP 9000

To boot a saved configuration on an HP 9000 system, interrupt the automatic boot process when you arrive at the boot console handler. Tell it to boot from the desired device (typically with a `boot pri` command). When it asks if you want to interact with the ISL or IPL, say Yes. (The exact mechanism to get to this point varies; consult your system's hardware document or the *hpxx(1M)* manpage for details.) At the **ISL>** prompt, type

```
ISL>hpxx thursday/vmunix
```

In either case, this will boot the saved configuration named `thursday`. When the boot is complete, it will be the currently running configuration; the previous configuration is lost (unless it was automatically saved as backup).

## Booting in Fail-Safe Mode

The other alternative for recovering from an unbootable configuration is to boot in fail-safe mode. When you boot the system in fail-safe mode, your configuration settings are ignored. All kernel tunables are given fail-safe values and no kernel modules are dynamically loaded during boot. This method is particularly useful when a hardware change or failure has caused all of your saved configurations to be unbootable.

## Integrity

To boot an Integrity system in fail-safe mode, get to the **HPUX>** prompt as described in “[Booting a Saved Configuration](#)” (page 198) and enter:

```
HPUX>boot -tm
```

## HP 9000

To boot an HP 9000 system in fail-safe mode, get to the **ISL>** prompt as described in “[Booting a Saved Configuration](#)” (page 198) and enter:

```
ISL>hpxx -tm
```

Some HP 9000 systems that have been updated from earlier versions of HP-UX have boot loaders that do not support the `-tm` option. On those systems, enter the following instead:

```
ISL>hpxx -f0x40000
```

(The two methods can be combined, if you want to boot a saved configuration in fail-safe mode. This uses the kernel executable built for the saved configuration, including all of its static modules, but none of its dynamically loaded modules.)

When you boot the system in fail-safe mode, the previous kernel configuration will be automatically saved for you, with a configuration name something like `saved_3DE78FA0`. The exact name will be printed for you in the boot messages on the console.

When you boot the system in fail-safe mode, the boot will stop when you reach single-user mode. At this time you should take any necessary steps to repair your system or your configuration and then reboot onto a valid configuration. HP does not recommend continuing to boot to multiuser mode after a fail-safe boot.

## Modifying Tunable Values at Boot Time

Tunable values can be changed at boot time by putting tunable settings at the end of the boot loader command line. These settings have the form `tunablename=value`, with no embedded white space. The value must be an integer in either hexadecimal (prefixed by `0x`) or decimal. This value will replace the value saved in the configuration being booted. (If the value is invalid, no change is made.)

For example, to change the value of `nproc` at boot time, do the following:

1. Boot the system and stop at the `HPUX>` (Integrity) or `ISL>` (HP 9000) prompt, as described in “Booting a Saved Configuration” (page 198).
2. Execute the `boot` (Integrity) or `hpx` (HP 9000) command with optional options and device file, followed by the tunable assignments, using the following syntax:

`boot [option]... [devicefile] [tunable=value]...`

`hpx [option]... [boot] [devicefile] [tunable=value]...`

For example, to boot the backup configuration and set `nproc` to 6000, the minimum commands would be:

`HPUX> boot backup nproc=6000`

`ISL> hpx backup/vmunix nproc=6000`

## Guidelines for Recovering from Errors

If you have an undesirable or unbootable kernel configuration, HP recommends the following approach to resolving the problem.

- If your system is up:
  - If you know which configuration change caused the problem:
    - If your backup configuration hasn't been updated since the bad change:
      - Load the backup configuration with `kconfig -l backup`.
    - Else (your backup configuration also has the problem in it):
      - Try to reverse the change using `kmodule` or `kctune`.  
Always specify the `-K` option to preserve the backup configuration.
  - Else (you don't know what change caused the problem, or the above didn't work):
    - Load a known good configuration using `kconfig -l`.  
Try the backup configuration first.
- Else (your system is down):
  - If you have had a hardware failure and now the system won't boot or if you need to preserve the bad configuration:
    - Try booting in fail-safe mode (see above).
    - Repair the configuration or the hardware, then reboot.
  - Else (no hardware failure, no need to preserve bad configuration):
    - Try booting a known good configuration, such as `backup`.

Of course, depending on the level of your support contract with HP, you can call on HP field service personnel to perform these steps, if needed.

If you get to a point where you cannot boot any of your saved configurations, even in fail-safe mode, your last resort is to boot from the HP-UX installation media. If that succeeds, you do not necessarily have to reinstall HP-UX; you can open a shell and try to repair your system.

## Kernel Configuration Example

In this example, the system administrator, Susan, is setting up a new HP-UX system to run a database server called `Prophet`. It has just finished booting after the initial install.

```
GenericSysName [HP Release B.11.31] (see /etc/issue)
Console Login: root
Password:
...
WARNING: YOU ARE SUPERUSER !!
```

The first thing Susan does is save a copy of the initial kernel configuration, in case she needs it later. She puts comments on all of her changes (with -C). She also puts a title on the saved configuration (with -t) to remind herself of what it contains.

```
kconfig -C "Save initial installation config" -s installed
 * The currently running configuration has been saved as
 'installed'.
kconfig -t installed "Initial installation"
 * The title of the configuration 'installed' has been set to
 "Initial installation".
```

The document for Prophet tells Susan to set the maxdsiz tunable to at least 0.5 TB, to set the semmni tunable to 3000, and to add 50 to whatever value she's using for shmmni. She also decides to add 1000 to the current value of nproc. Susan starts by looking at the current values of these tunables, and the descriptions of the ones she's unfamiliar with.

```
kctune nproc maxdsiz
Tunable Value Expression Changes
maxdsiz 1073741824 Default Immed
nproc 4200 Default Immed
kctune -d semmni shmmni
Tunable Value Expression Changes
Description
semmni 2048 Default
 Maximum number of semaphore sets on the system
shmmni 400 Default Immed
 Maximum number of shared memory segments on the system
```

Having done that, she sets the values as directed. She sets them all on the same command line so that they will all take effect at the same time. Since one of the changes cannot be made immediately, all of the changes are held for next boot.

```
kctune -C "Tunable settings for Prophet" "nproc+=1000" "maxdsiz>=512000000" \
 "semmni=3000" "shmmni+=50"
NOTE: The requested changes could not be applied to the currently
 running system, for the following reasons:
 - The tunable 'semmni' cannot be changed without a reboot.
 * The requested changes have been saved, and will take effect at
 next boot.
NOTE: No change to the tunable 'maxdsiz' was needed.
Tunable Value Expression Changes
maxdsiz 1073741824 Default Immed
nproc (now) 4200 Default Immed
 (next boot) 5200 5200
semmni (now) 2048 Default
 (next boot) 3000 3000
shmmni (now) 400 Default Immed
 (next boot) 450 450
```

Susan also decides to remove an unused LAN driver. First, she verifies which drivers are currently in use.

```
nwmgr
```

| Name/<br>Class | Interface<br>Instance | Station<br>State | Sub-<br>system | Interface<br>Type | Related<br>Interface |
|----------------|-----------------------|------------------|----------------|-------------------|----------------------|
|----------------|-----------------------|------------------|----------------|-------------------|----------------------|

```
===== ===== ===== ===== ===== ===== =====
lan0 UP 0x00306E4949FD gelan 1000Base-T
```

Then she verifies that the modules are installed.

```
kcmodule -d gelan iether
Module State Cause Notes
Description
gelan static best loadable, unloadable
 Gigabit Ethernet (gelan) LAN Driver
iether static best loadable, unloadable
 Intel 8254X Ethernet Driver (for 100BT and Gigabit Cards)
```

Finally, she removes the iether driver.

```
kcmodule -C "removing unneeded iether Ethernet driver" iether=unused
Building a new kernel for the configuration to be used at next boot... done.
NOTE: The requested changes could not be applied to the currently
running system, for the following reasons:
- Moving a module into or out of the 'static' state requires a
kernel rebuild.
* The requested changes have been saved, and will take effect at
next boot.
Module State Cause Notes
iether (now) static best loadable, unloadable
 (next boot) unused
```

Since iether is static, a new kernel is built, and marked for use at next boot.

Susan checks a summary of all of her changes that will take effect when she reboots.

```
kconfig -D
Module State Cause Notes
iether (now) static best loadable, unloadable
 (next boot) unused
Tunable Value Expression Changes
nproc (now) 4200 Default Immed
 (next boot) 5200 5200
semmni (now) 2048 Default
 (next boot) 3000 3000
shmmni (now) 400 Default Immed
 (next boot) 450 450
```

Satisfied, she reboots.

```
shutdown -r
...
GenericSysName [HP Release B.11.31] (see /etc/issue)
Console Login: root
Password:
...
WARNING: YOU ARE SUPERUSER !!
```

After the reboot, Susan verifies the configuration status.

```
kconfig
Configuration Title
backup Automatic Backup
```

```

installed Initial installation
last_install Created by last OS install
kconfig -w
 * The currently running configuration was created on Fri Dec 15
 17:39:39 2006 by root as a copy of 'last_install'.
 * It was last saved on Fri Dec 15 17:45:11 2006 by root.
 * It was last modified on Fri Dec 15 18:04:58 2006 by root.

```

Then, she saves the new kernel configuration under the name good, so that she can go back to it, if needed. She gives it a title to help recognize it later.

```

kconfig -C "Good configuration for Prophet" -s good
 * The currently running configuration has been saved as 'good'.
kconfig -t good "Good configuration for Prophet"
 * The title of the configuration 'good' has been set to "Good
 configuration for Prophet".

```

After some time, one of her users asks her to increase the size of the buffer cache, hoping to speed up the application. She complies — after all, it doesn't need a reboot, so she can do it without disturbing anyone. Since it's the first change after a boot, the system asks whether to make automatic backups.

```

kctune -C "Bigger file cache for better performance" filecache_max=20%
==> Update the automatic 'backup' configuration first? yes
 * The automatic 'backup' configuration has been updated.
 * Future operations will update the backup without prompting.
 * The requested changes have been applied to the currently
 running configuration.
Tunable Value Expression Changes
filecache_max (before) 1017118720 Default Imm (auto disabled)
 (now) 406847488 20%

```

It's a good thing she said yes. The larger buffer cache actually slowed things down — but all she has to do is restore the automatic backup.

```

kconfig -C "Putting file cache back; performance was worse." -l backup
 * The automatic 'backup' configuration has been updated.
 * The requested changes have been applied to the currently
 running configuration.
 * The automatic 'backup' configuration has been loaded and is now
 in use.
kctune filecache_max
Tunable Value Expression Changes
filecache_max 1017118720 Default Auto

```

While Susan is on vacation, her colleague, Fred, decides to use the machine for billing software during the night. This software needs to execute code on the stack (a security risk), so he enables that behavior (which is prohibited by default). No reboot is needed to do so.

```

kctune -d executable_stack
Tunable Value Expression Changes
Description
executable_stack 0 Default Immed
 Enables execution of code on a stack (0 = no, 1 = yes, 2 = yes but warn)

```

```
kctune -C "Nightly billing s/w needs execute-on-stack" executable_stack=1
 * The automatic 'backup' configuration has been updated.
 * The requested changes have been applied to the currently
 running configuration.
Tunable Value Expression Changes
executable_stack (before) 0 Default Immed
 (now) 1
```

The billing software also uses the kernel Random Number Generator module. Fred checks and sees that it's not in use, but since it's loadable he doesn't need to reboot to use it.

```
kcmodule -d rng
Module State Cause Notes
 Description
rng unused loadable, unloadable
 Strong Random Number Generator
```

He goes ahead and loads the module.

```
kcmodule -C "Random Number Generator needed for nightly billing jobs" rng=best
 * The automatic 'backup' configuration has been updated.
 * The requested changes have been applied to the currently
 running configuration.
Module State Cause Notes
rng (before) unused loadable, unloadable
 (now) loaded best
 (next boot) loaded explicit
```

Fred saves these new configuration settings under the name night, with a descriptive title.

```
kconfig -C "Settings for nightly billing jobs" -s night
 * The currently running configuration has been saved as 'night'.
kconfig -t night "Nightly billing jobs"
 * The title of the configuration 'night' has been set to "Nightly
 billing jobs".
```

Since good isn't a very helpful name for Susan's configuration anymore, Fred renames it to day. He checks the list of configurations to make sure everything looks OK.

```
kconfig -r good day
 * The configuration 'good' has been renamed to 'day'.

kconfig
Configuration Title
backup Automatic Backup
day Good configuration for Prophet
installed Initial installation
last_install Created by last OS install
night Nightly billing jobs
```

Finally, he tries loading first the day configuration, and then the night configuration, to make sure he can move back and forth at will.

```
kconfig -l day
 * The automatic 'backup' configuration has been updated.
 * The requested changes have been applied to the currently
```

```

 running configuration.
* The configuration 'day' has been loaded and is now in use.

kconfig -l night
 * The automatic 'backup' configuration has been updated.
 * The requested changes have been applied to the currently
 running configuration.
 * The configuration 'night' has been loaded and is now in use.

When Susan returns from her vacation, the first thing she does is check the automatically
maintained log file to see what Fred has done.

kclog 5
=====
2006-12-15 18:28:45 MST root:
kconfig -C 'Settings for nightly billing jobs' -s night

 * The currently running configuration has been saved as 'night'.

=====
2006-12-15 18:29:07 MST root:
kconfig -t night 'Nightly billing jobs'

 * The title of the configuration 'night' has been set to "Nightly
 billing jobs".

=====
2006-12-15 18:30:07 MST root:
kconfig -r good day

 * The configuration 'good' has been renamed to 'day'.

=====
2006-12-15 18:30:55 MST root:
kconfig -l day

 * The configuration 'day' has been loaded and is now in use.

=====
2006-12-15 18:31:20 MST root:
kconfig -l night

 * The configuration 'night' has been loaded and is now in use.

```

She can see that Fred has put a new application on her server, and worse, an insecure one. At least he tested and documented his changes.

Susan doesn't want to leave her system the way Fred changed it, so she moves the nightly billing job to another system. First, she exports his night configuration to a text file.

```
kconfig -e night /tmp/system.night
 * The configuration 'night' has been exported to
 '/tmp/system.night'.
```

Moving the file over to another machine, she imports the configuration there, using the -V option to ensure that exactly the same kernel software is in use. Then she loads the configuration. Something about the configuration can't be changed immediately — probably a tunable setting — so she has to reboot the machine. As intended, the machine uses Fred's night configuration when it comes back up.

```
kconfig -C "Move nightly billing jobs here from Prophet" -iv night \
/tmp/system.night
 * The configuration 'night' has been imported from
 '/tmp/system.night'.

kconfig -l night
ERROR: The requested changes could not be applied to the currently
running system, for the following reasons:
- Moving a module into or out of the 'static' state requires a
 kernel rebuild.
* The configuration 'night' has been marked for use at next boot.

shutdown -r
```

## Kernel Configuration Quick Reference Tables

**Table 10-11 Working with Kernel Configurations**

| Procedure                                | Command                       |
|------------------------------------------|-------------------------------|
| Choose the configuration to boot...      |                               |
| ...before the reboot <sup>1</sup>        | kconfig [-f] -n configname    |
| ...at the boot loader prompt (Integrity) | boot configname               |
| ...at the boot loader prompt (HP 9000)   | hpx configname/vmunix         |
| List all kernel configurations           | kconfig [-v]                  |
| Save the currently running configuration | kconfig [-f] -s newname       |
| Copy a saved configuration               | kconfig -c src dest           |
| Rename a saved configuration             | kconfig -r old new            |
| Delete a saved configuration             | kconfig [-f] -d configname    |
| Load a saved configuration               | kconfig [-f] -l configname    |
| Set the title of a configuration         | kconfig -t configname "title" |

1 If this option is used, there is no need to interrupt the boot process to select the new kernel configuration.

**Table 10-12 Working with System Files**

| Procedure                                                         | Command                                     |
|-------------------------------------------------------------------|---------------------------------------------|
| Create a system file...                                           |                                             |
| ...for a saved configuration                                      | <code>kconfig -e configname filename</code> |
| ...for the currently running configuration <sup>1</sup>           | <code>kconfig -e filename</code>            |
| Create/update a configuration from a system file ... <sup>2</sup> |                                             |
| ... create/update a saved configuration                           | <code>kconfig -i configname filename</code> |
| ...update the currently running configuration                     | <code>kconfig [-fhV] -i filename</code>     |

1 Includes any changes being held for next boot.

2 `mk_kernel` can also be used for this purpose.

**Table 10-13 Working with Changes Held for Next Boot**

| Procedure                                                                                                                                                                                           | Command                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>NOTE:</b> <code>kconfig -i</code> , <code>kcmodule</code> , and <code>kctune</code> hold their changes until next boot if they can't be applied immediately, or if <code>-h</code> is specified. |                         |
| List all changes being held for next boot                                                                                                                                                           | <code>kconfig -D</code> |
| Discard all changes being held for next boot                                                                                                                                                        | <code>kconfig -H</code> |

**Table 10-14 Working with Tunables**

| Procedure                                         | Command                                  |
|---------------------------------------------------|------------------------------------------|
| List tunables and their values...                 | <code>kctune [tunable]...</code>         |
| ...verbose output                                 | <code>-v</code>                          |
| ...only tunables with changes held for next boot  | <code>-D</code>                          |
| ...include derived tunables set to default values | <code>-a</code>                          |
| ...group by module name                           | <code>-g</code>                          |
| ...in a saved configuration                       | <code>-c configname</code>               |
| Set a tunable value                               | <code>kctune tunable="expression"</code> |
| Set a tunable to default                          | <code>kctune tunable=default</code>      |
| Increment a tunable value                         | <code>kctune tunable+=value</code>       |
| Make sure tunable value is at least n             | <code>kctune "tunable&gt;=n"</code>      |
| ...hold change until next boot                    | <code>-h</code>                          |

**Table 10-14 Working with Tunables (continued)**

| <b>Procedure</b>                       | <b>Command</b>             |
|----------------------------------------|----------------------------|
| ...apply change to saved configuration | <code>-c configname</code> |
| ...create user-defined tunable         | <code>-u</code>            |

**Table 10-15 Working with Kernel Modules**

| <b>Procedure</b>                                | <b>Command</b>                      |
|-------------------------------------------------|-------------------------------------|
| List modules and their states...                | <code>kcmodule [module]...</code>   |
| ...verbose output                               | <code>-v</code>                     |
| ...only modules with changes held for next boot | <code>-D</code>                     |
| ...include required modules                     | <code>-a</code>                     |
| ...in a saved configuration                     | <code>-c configname</code>          |
| Add a module to the configuration...            |                                     |
| ...in default state                             | <code>kcmodule module=best</code>   |
| ...statically bound into the kernel executable  | <code>kcmodule module=static</code> |
| ...dynamically loaded, now and at each boot     | <code>kcmodule module=loaded</code> |
| ...auto-loaded at first use                     | <code>kcmodule module=auto</code>   |
| Remove a module from the configuration...       | <code>kcmodule module=unused</code> |
| ...Hold change until next boot                  | <code>-h</code>                     |
| ...Apply change to saved configuration          | <code>-c configname</code>          |

**Table 10-16 Working with the Kernel Configuration Log File**

| <b>Procedure</b>                                                                                                                       | <b>Command</b>                                      |
|----------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| <b>NOTE:</b> The log file is located at <code>/var/adm/kc.log</code> . The <code>kc*</code> commands add a log entry for every change. |                                                     |
| Add a comment to the log file...                                                                                                       |                                                     |
| ...while making a change with a <code>kc*</code> command                                                                               | <code>add -C "comment" to the change command</code> |
| ...without making a configuration change                                                                                               | <code>kclog -C "comment"</code>                     |
| View the last <i>n</i> entries in the log (default is 1)...                                                                            | <code>kclog n</code>                                |
| ...counting only changes to a configuration                                                                                            | <code>-c configname</code>                          |
| ...counting only changes of a particular type                                                                                          | <code>-t module tunable device</code>               |

**Table 10-16 Working with the Kernel Configuration Log File (continued)**

| Procedure                                        | Command                                         |
|--------------------------------------------------|-------------------------------------------------|
| ...counting only changes to a particular item    | <code>-n modulename tunablename hwpather</code> |
| ...counting only log entries containing a string | <code>-f "string"</code>                        |

**Table 10-17 Kernel Configuration File Locations**

| Procedure                                | Command                               |
|------------------------------------------|---------------------------------------|
| Saved configurations are stored in...    | <code>/stand/configname</code>        |
| Kernel executable is at...               | <code>/stand/configname/vmunix</code> |
| System file is at...                     | <code>/stand/configname/system</code> |
| Currently running configuration is in... | <code>/stand/current</code>           |
| Kernel executable is at...               | <code>/stand/current/vmunix</code>    |
| System file is at...                     | <code>/stand/current/system</code>    |

**NOTE:** Never directly manipulate any of the files in a kernel configuration directory, except the system file. Always use the `kc*` commands.

## Transition from Previous HP-UX Releases

Experienced administrators of releases prior to HP-UX 11i v2 will find some aspects of the kernel configuration mechanisms to be unfamiliar. However, many of the underlying concepts are unchanged. The tables in this section give information to help administrators to make the transition.

**Table 10-18 Kernel Configuration Methodology**

| HP-UX 11i Version 1 and Before                                                                                             | HP-UX 11i Version 2 and After                                                                                                                                                                       |
|----------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Use SAM to configure the kernel.                                                                                           | Use HP SMH to configure the kernel.                                                                                                                                                                 |
| Look at <code>/stand/system</code> to see the current configuration.                                                       | Same. <sup>1</sup>                                                                                                                                                                                  |
| Run an unsupported command to make sure <code>/stand/system</code> is up to date.                                          | Not needed. <sup>1</sup> <code>/stand/system</code> is automatically kept up to date.                                                                                                               |
| Make configuration changes by editing <code>/stand/system</code> and running <code>mk_kernel</code> .                      | Same. Changes will be applied to the running system (no reboot), if possible. <sup>1</sup>                                                                                                          |
| Make configuration changes by running <code>kmtune</code> or <code>kmsystem</code> , then running <code>mk_kernel</code> . | Make the changes with <code>kctune</code> or <code>kcmodule</code> (no <code>mk_kernel</code> ), or edit <code>/stand/system</code> manually <sup>1 2 3</sup> and then run <code>mk_kernel</code> . |

**Table 10-18 Kernel Configuration Methodology (continued)**

| <b>HP-UX 11i Version 1 and Before</b>                                                        | <b>HP-UX 11i Version 2 and After</b>     |
|----------------------------------------------------------------------------------------------|------------------------------------------|
| Make configuration changes by editing /stand/system and running config.                      | Use mk_kernel instead. <sup>1</sup>      |
| Manage DLKMs with the kminstall, kmsystem, kmmodreg, kmadmin, kmupdate, and config commands. | Manage DLKMs using kmodule. <sup>2</sup> |
| View or change tunables using kmtune. <sup>4</sup>                                           | Use kctune instead. <sup>3</sup>         |

1 See “Managing Configurations with System Files” (page 192).

2 See “Managing Kernel Modules with kmodule” (page 160).

3 See “Managing Kernel Tunable Parameters with kctune” (page 170).

4 HP-UX 11i v2 contained a compatibility stub for kmtune. It was removed in HP-UX 11i v3.

**Table 10-19 Kernel Configuration Commands and Options**

| <b>HP-UX 11i Version 1 and Before</b> | <b>HP-UX 11i Version 2 and After</b>   |
|---------------------------------------|----------------------------------------|
| config (without -M)                   | mk_kernel <sup>1</sup>                 |
| config -M                             | No longer needed                       |
| kmadmin -b                            | No longer needed                       |
| kmadmin -k                            | kmodule <sup>2</sup>                   |
| kmadmin -L modulename                 | kmodule modulename=loaded <sup>2</sup> |
| kmadmin -U modulename                 | kmodule modulename=unused <sup>2</sup> |
| kmadmin -u module_id                  | kmodule modulename=unused <sup>2</sup> |
| kmadmin -q module_id                  | kmodule -v modulename <sup>2</sup>     |
| kmadmin -Q modulename                 | kmodule -v modulename <sup>2</sup>     |
| kmadmin -s                            | kmodule <sup>2</sup>                   |
| kmadmin -S                            | kmodule -v <sup>2</sup>                |
| kminstall                             | No longer needed                       |
| kmmodreg                              | No longer needed                       |
| kmpath (no options) <sup>3</sup>      | kcpath -x                              |
| kmpath -k <sup>3</sup>                | kcpath -b                              |
| kmpath -c <sup>3</sup>                | kcpath -d                              |
| kmpath -i <sup>3</sup>                | No longer needed                       |
| kmsystem (no options)                 | kmodule <sup>2</sup>                   |

**Table 10-19 Kernel Configuration Commands and Options (continued)**

| <b>HP-UX 11i Version 1 and Before</b>                | <b>HP-UX 11i Version 2 and After</b>                 |
|------------------------------------------------------|------------------------------------------------------|
| <code>kmsystem -b</code>                             | No longer needed                                     |
| <code>kmsystem -c y -l y modulename</code>           | <code>kcmodule modulename=loaded</code> <sup>2</sup> |
| <code>kmsystem -c y -l n modulename</code>           | <code>kcmodule modulename=static</code> <sup>2</sup> |
| <code>kmsystem -c n modulename</code>                | <code>kcmodule modulename=unused</code> <sup>2</sup> |
| <code>kmsystem -q modulename</code>                  | <code>kcmodule -v modulename</code> <sup>2</sup>     |
| <code>kmtune (no options)</code> <sup>3</sup>        | <code>kctune</code> <sup>4</sup>                     |
| <code>kmtune -l</code> <sup>3</sup>                  | <code>kctune -v</code> <sup>4</sup>                  |
| <code>kmtune -q tunable</code> <sup>3</sup>          | <code>kctune tunable</code> <sup>4</sup>             |
| <code>kmtune -r tunable</code> <sup>3</sup>          | <code>kctune tunable=Default</code> <sup>4</sup>     |
| <code>kmtune -u -s tunable=value</code> <sup>3</sup> | <code>kctune tunable=value</code> <sup>4</sup>       |
| <code>kmtune -u -s tunable+value</code> <sup>3</sup> | <code>kctune tunable+=value</code> <sup>4</sup>      |
| <code>kmtune -s tunable=value</code> <sup>3</sup>    | <code>kctune -h tunable=value</code> <sup>4</sup>    |
| <code>kmupdate (no options)</code>                   | <code>kconfig -n hpxx_test</code> <sup>5</sup>       |
| <code>kmupdate kernel</code>                         | <code>kconfig -n configuration</code> <sup>5</sup>   |
| <code>kmupdate -M module</code>                      | No longer needed                                     |
| <code>kmupdate -d kernel</code>                      | <code>kconfig -d configuration</code> <sup>6</sup>   |
| <code>mk_kernel (without -M)</code>                  | <code>mk_kernel</code> <sup>1</sup>                  |
| <code>mk_kernel -M</code>                            | No longer needed                                     |

1 See “Managing Configurations with System Files” (page 192).

2 See “Managing Kernel Modules with kcmodule” (page 160).

3 HP-UX 11i v2 contained compatibility stubs for `kmpath` and `kmtune`. They were removed in HP-UX 11i v3.

4 See “Managing Kernel Tunable Parameters with kctune” (page 170).

5 See “Using Saved Configurations” (page 191).

6 See “Modifying Saved Configurations” (page 192).

**Table 10-20 Kernel Configuration Files and Directories**

| <b>HP-UX 11i Version 1 and Before</b>                | <b>HP-UX 11i Version 2 and After</b>                   |
|------------------------------------------------------|--------------------------------------------------------|
| Currently running kernel: <code>/stand/vmunix</code> | <code>/stand/vmunix</code>                             |
| Backup kernel: <code>/stand/vmunix.prev</code>       | Backup configuration: <code>backup</code> <sup>1</sup> |

**Table 10-20 Kernel Configuration Files and Directories (continued)**

| HP-UX 11i Version 1 and Before                                                       | HP-UX 11i Version 2 and After                                                                                                                                          |
|--------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Test kernel: /stand/build/vmunix_test<br>(default output of <code>mk_kernel</code> ) | Test configuration: <code>hpxx_test</code> <sup>2</sup>                                                                                                                |
| Primary system file: /stand/system                                                   | /stand/system <sup>2</sup>                                                                                                                                             |
| Module system files: /stand/system.d/*                                               | No longer used. The data are now in the primary system file, /stand/system. <sup>2</sup>                                                                               |
| Master files: /usr/conf/master.d/*                                                   | No longer used. The data are embedded into the kernel code, and available through the <code>kmodule</code> <sup>3</sup> and <code>kctune</code> commands. <sup>4</sup> |

1 See “The Automatic Backup Configuration” (page 198).

2 See “Managing Configurations with System Files” (page 192).

3 See “Managing Kernel Modules with `kmodule`” (page 160).

4 See “Managing Kernel Tunable Parameters with `kctune`” (page 170).



---

# Index

## Symbols

\$HOME/.cshrc file, 42, 68  
\$HOME/.forward file, 54  
\$HOME/.kshrc file, 54  
\$HOME/.login file, 42, 68  
\$HOME/.profile file, 42, 68  
\$HOME/.rhosts file, 54  
    example, 72  
/dev/ttyconf file, 37  
/etc/csh.login file, 42, 68  
/etc/default/autofs file, 80  
/etc/default/security file, 36  
/etc/dfs/dfstab file, 72, 84, 85  
/etc/exports file, 75  
/etc/fstab file, 80, 82, 84, 85  
/etc/group file, 67, 88  
/etc/hosts file, 84, 85  
/etc/hosts.equiv file, 72  
/etc/inetd.conf file, 90  
/etc/inittab file, 37, 38  
/etc/loggingroup file, 67  
/etc/mail/aliases file, 122  
/etc/nologin file, 36  
/etc/passwd file, 47, 50, 62, 67, 87, 88  
/etc/profile file, 42, 68  
/etc/rc.config.d/mailservs file, 125  
/etc/rc.config.d/nfsconf file, 84  
/etc/resolv.conf file, 84, 85  
/etc/services file, 114  
/etc/shadow file, 50  
/etc/shells file, 90  
/etc/skel directory, 54  
/etc/skel/.cshrc file, 54, 68  
/etc/skel/.exrc file, 54  
/etc/skel/.login file, 54, 68  
/etc/skel/.profile file, 54, 68  
/opt/hpnp/admin/jetadmin command, 116  
/opt/hpsmh/bin/hpsmh command, 25  
/sbin/init command, 38  
/sbin/mount command, 72  
/sbin/set\_parms command, 40, 41  
/sbin/shutdown command, 38  
/sbin/stty command, 37  
/usr/bin/cancel command, 118  
/usr/bin/chacl command, 68  
/usr/bin/chgrp command, 68  
/usr/bin/chmod command, 68  
/usr/bin/chown command, 62, 68  
/usr/bin/elm command, 121  
/usr/bin/getacl command, 68  
/usr/bin/lpstat command, 112, 114, 118, 120  
/usr/bin/lsacl command, 68

/usr/bin/mail command, 121  
/usr/bin/mkdir command, 62  
/usr/bin/newalias command, 122  
/usr/bin/passwd command, 62  
/usr/bin/setacl command, 68  
/usr/bin/umask command, 68  
/usr/bin/uux command, 121  
/usr/bin/X11/xhost command, 73  
/usr/sbin/accept command, 112, 114, 117  
/usr/sbin/automount command, 80  
/usr/sbin/automountd daemon, 80  
/usr/sbin/ch\_rc command, 36  
/usr/sbin/disable command, 118  
/usr/sbin/enable command, 112, 114  
/usr/sbin/exportfs command, 75  
/usr/sbin/fsweb command, 30  
/usr/sbin/getty command, 37  
/usr/sbin/inetd daemon, 90  
/usr/sbin/kcweb command, 30  
/usr/sbin/ladmin command, 112, 114, 117, 118, 119, 120  
/usr/sbin/lpmove command, 118, 120  
/usr/sbin/lpsched command, 112, 114, 117, 118, 119, 120  
/usr/sbin/lpshut command, 112  
/usr/sbin/mount command, 72, 82  
/usr/sbin/ncweb command, 30  
/usr/sbin/pdweb command, 30  
/usr/sbin/reject command, 118, 120  
/usr/sbin/rlpdaemon command, 114  
/usr/sbin/secweb command, 30  
/usr/sbin/sendmail command, 121, 125  
/usr/sbin/setprivgrp command, 67  
/usr/sbin/share command, 72, 75  
/usr/sbin/shareall command, 85  
/usr/sbin/smh command, 30  
/usr/sbin/stty command, 37  
/usr/sbin/ugweb command, 30  
/usr/sbin/unshare command, 85  
/usr/sbin/useradd command, 54, 60  
/usr/sbin/userdel command, 60  
/usr/sbin/usermod command, 60  
/usr/sbin/vipw command, 47, 62  
/var/adm/inetd.sec file, 114  
/var/spool/mqueue file, 122

## A

access control lists (*see* ACLs)  
account aging  
    shadow field, 51  
accounts  
    aging, 51  
ACLS  
    file protection, 68  
    HFS file system, 68

JFS file system, 68  
 add a group  
   HP SMH, text-based, 59  
 add a user  
   command, 60  
   commands, 61  
   HP SMH, text-based, 50  
     with a template, 55  
   useradd command, 60  
 AUTOMOUNT\_TIMEOUT environment variable, 80

**C**

case studies  
   user on multiple systems, 70  
 characters  
   eof, 37  
   erase, 37  
   intr, 37  
   kill, 37  
   quit, 37  
   start, 37  
   stop, 37  
 CIFS  
   sharing directories between HP-UX and Windows, 107  
 clients  
   X Window System, 73  
 commands  
   accept, 112, 114, 117  
   add a user, 60, 61  
   adding local printer to spooler, 112  
   adding network printer to spooler, 116  
   adding printer class to spooler, 117  
   adding remote printer to spooler, 114  
   automount, 80  
   cancel, 118  
   ch\_rc, 36  
   chacl, 68  
   chgrp, 68  
   chmod, 68  
   chown, 62, 68  
   configuring anonymous FTP, 88  
   disable, 118  
   elm, 121  
   enable, 112, 114  
   exportfs, 75  
   fsweb, 30  
   getacl, 68  
   getty, 37  
   hpsmh, 25  
   init, 38  
   jetadmin, 116  
   kcweb, 30  
   lpadmin, 112, 114, 117, 118, 119, 120  
   lpmove, 118, 120  
   lpsched, 112, 114, 117, 118, 119, 120  
   lpshut, 112  
   lpstart, 112  
   lpstat, 112, 114, 118, 120  
   lsacl, 68  
   mail, 121  
   mailx, 121  
   mixing with HP SMH, 20  
   mkdir, 62  
   mount, 72, 82  
   ncweb, 30  
   newalias, 122  
   NFS client  
     automount shared file system, 82  
   NFS server  
     share directory, 77  
   passwd, 62  
   pdweb, 30  
   reject, 118, 120  
   removing printer class, 120  
   removing printer from class, 119  
   removing printer from spooler, 118  
   rlpdaemon, 114  
   secweb, 30  
   sendmail, 121, 125  
   set\_parms, 40, 41  
   setacl, 68  
   setprivgrp, 67  
   share, 72, 75  
   shareall, 85  
   shutdown, 38  
   smh, 30  
   smh -r, 31  
   stty, 37  
   summary, 20  
   syntax, 16  
   ugweb, 30  
   umask, 68  
   unshare, 85  
   useradd, 54, 60  
   userdel, 60  
   usermod, 60  
   uux, 121  
   versus HP SIM and HP SMH, 20  
   vipw, 47, 62  
   xhost, 73  
 comments  
   passwd field, 51  
 Common Internet File System (*see CIFS*)  
 configuration tools, 19  
 conventions  
   shells, 16  
 cross-system access  
   user ID, 69

**D**

daemons  
   /usr/sbin/inetd, 90  
   automountd, 80  
 date

problems, 40  
set\_parms command, 40  
setting, 40  
setting with set\_parms command, 42

delete a user  
HP SMH, text-based, 58  
userdel command, 60

directories  
/etc/skel, 54  
/homeroot, 47  
home, 50  
local, 70  
remote, 70  
case study, 70  
skeleton, 54

DISPLAY environment variable, 30

DNS server  
setting with set\_parms command, 42

domain name  
setting with set\_parms command, 42

Domain Name Service (*see DNS*)

## E

e-mail (*see mail*)

electronic mail (*see mail*)

environment variables, 11  
(*see also variables*)

AUTOMOUNT\_TIMEOUT, 80  
DISPLAY, 30  
NFS\_CLIENT, 84  
NFS\_CORE, 84  
NFS\_SERVER, 84  
SENDMAIL\_FREEZE, 125, 126, 127, 130  
SENDMAIL\_SERVER, 125, 126, 127, 130  
SENDMAIL\_SERVER\_NAME, 125, 126, 127, 130

eof character, 37

erase character, 37

error recovery

PCI, 133

Ethernet cards

assigning IP addresses and masks with HP SMH, web-based, 137  
configuring LAN, 139  
frequently asked questions, 137  
gigabit Ethernet, 141  
installing network driver, 136  
jumbo frames, 141  
related information, 142  
routing, 139  
setting speed and duplex, 140  
verifying driver in kernel, 142  
verifying LAN installation, 138  
verifying speed and duplex, 140

examples

user on multiple systems, 70

## F

file access, 67

file protection  
ACLs, 68

file systems  
remote  
case study, 70  
mounting, 70  
sharing, 70

File Transfer Protocol (*see FTP*)

files

\$HOME/.cshrc, 42, 68  
\$HOME/.forward, 54  
\$HOME/.kshrc, 54  
\$HOME/.login, 42, 68  
\$HOME/.profile, 42, 68  
\$HOME/.rhosts, 54  
example, 72  
/dev/ttyconf, 37  
/etc/csh.login, 42, 68  
/etc/default/autofs, 80  
/etc/default/security, 36  
/etc/dfs/dfstab, 72, 75, 84, 85  
/etc/exports, 75  
/etc/fstab, 80, 82, 84, 85  
/etc/group, 67, 88  
/etc/hosts, 84, 85  
/etc/hosts.equiv, 72  
/etc/inetd.conf, 90  
/etc/inittab, 37, 38  
/etc/logginggroup, 67  
/etc/mail/aliases, 122  
/etc/nologin, 36  
/etc/passwd, 47, 50, 62, 67, 87, 88  
/etc/profile, 42, 68  
/etc/rc.config.d/mailservs, 125  
/etc/rc.config.d/nfsconf, 84  
/etc/resolv.conf, 84, 85  
/etc/services, 114  
/etc/shadow, 50  
/etc/shells, 90  
/etc/skel/.cshrc, 54, 68  
/etc/skel/.exrc, 54  
/etc/skel/.login, 54, 68  
/etc/skel/.profile, 54, 68  
/var/adm/inetd.sec, 114  
/var/spool/Imqueue, 122  
access permissions, 68  
ownership, 68

format

functions, 16

FTP

anonymous

/etc/group file, 88  
/etc/passwd file, 87, 88

access for users without logins, 87  
configuring with commands, 88

configuring with HP SMH, text-based, 87  
copy files between systems, 87  
login  
  /etc/inetd.conf file, 90  
  /etc/shells file, 90  
  /usr/sbin/inetd daemon, 90  
  configuring, 91  
  configuring with Windows, 91  
  connecting from HP-UX to Windows, 91  
  connecting from Windows to HP-UX, 99  
  getting file to HP-UX from Windows, 95  
  getting file to Windows from HP-UX, 102  
  getusershell() function, 90  
  sending file from HP-UX to Windows, 96  
  sending file to HP-UX from Windows, 104  
  test with HP SMH, text-based, 90  
  test with HP SMH, web-based, 90  
  troubleshooting, 90  
functions  
  definition format, 16  
  getusershell(), 90  
  syntax, 16  
  usage format, 16

**G**

gateway  
  setting with set\_parms command, 42

gecos  
  passwd field, 51

General Electric Comprehensive Operating Supervisor, 51

getusershell() function, 90

groups  
  add  
    HP SMH, text-based, 59  
  primary, 50, 67  
  secondary, 67  
  special privileges, 67

**H**

HFS file system  
  ACLs, 68

home directory  
  creating, 53  
  local, 70  
  mounting, 72  
  passwd field, 50  
  remote, 70  
    case study, 70  
  sharing, 72

home phone  
  passwd field, 51

host name  
  setting with set\_parms command, 42

HP Serviceguard  
  PCI error recovery, 133

HP SIM

summary, 19  
versus HP-UX commands, 20

HP SMH  
  users and groups, 49  
  versus HP-UX commands, 20

HP SMH, text-based  
  add a group, 59  
  add a user, 50  
    with a template, 55  
  adding local printer to spooler, 112  
  adding remote printer to spooler, 113  
  configuring anonymous FTP, 87  
  create user template, 55  
  delete a user, 58  
  functional areas, 25  
  modify a user, 57  
  modules, 30  
  NFS client  
    automount shared file system, 82  
    enable or disable, 76  
    standard-mount shared file system, 80

NFS server  
  enable or disable, 76  
  share directory, 77

restricted access, 31

starting, 30

summary, 20

test FTP login, 90

X Window, 30

HP SMH, web-based  
  adding network printer to spooler, 116  
  daemon  
    target system, 25

Ethernet cards  
  assigning IP addresses and masks, 137

functional areas, 21

managing kernel, 165

starting, 24

summary, 20

test FTP login, 90

web client, 24

HP Systems Insight Manager (*see* HP SIM)

HP-UX  
  operating system version, 13, 14  
  release name, 13

HP-UX commands (*see* commands)

HP-UX release information, 12

HP-UX System Management Homepage (*see* HP SMH)

|

intr character, 37

IP address  
  setting with set\_parms command, 42

J

JFS file system

ACLs, 68

## K

### kernel

- automatic backup, 198
- booting in fail-safe mode, 199
- booting saved configuration, 198
- configuration command overview, 154
- configuration commands
  - common behavior, 157
- configuration example, 201
- configuration features, 153
- configuration log file, 195
- configuration purpose, 154
- configuration with HP SMH, 155
- error recovery, 197
- error recovery guidelines, 201
- kcmodule command, 160
- managing configurations with system files, 192
- managing running configuration with kconfig command, 188
- managing saved configuration with kconfig command, 189
- managing with HP SMH, 165
- modifying tunables at boot time, 200
- monitoring alarms
  - activate, deactivate, remove, 187
- monitoring alarms with HP SMH, web-based, 180
- monitoring resource usage with HP SMH, web-based, 180
- other configuration operations, 156
- parsable command output, 196
- quick reference tables, 207
- transition from previous HP-UX releases, 210
- uses for system files, 194
- using and modifying saved configuration with kconfig command, 191

### kernel tunables

- managing with HP SMH, web-based, 175
- managing with ktune command, 170
- pci\_eh\_enable, 134
- pci\_error\_tolerance\_time, 134

### kill character, 37

## L

### line printer (*see* printers)

### line printer spooler (*see* spooler)

### location

- passwd field, 51

### login environments

- \$HOME/.cshrc, 68
- \$HOME/.login, 68
- \$HOME/.profile, 68
- /etc/csh.login, 68
- /etc/profile, 68
- /etc/skel/.cshrc, 68
- /etc/skel/.login, 68
- /etc/skel/.profile, 68

customizing, 68

system-wide and user, 42

### login files

\$HOME/.rhosts, 54

### login name

passwd field, 50

shadow field, 50

### login shell

passwd field, 51

login special characters, 37

LP spooler (*see* spooler)

## M

### mail

/etc/rc.config.d/mailservs file, 125

/usr/sbin/sendmail command, 125

/var/spool/mqueue file, 122

alias files, 122

aliases file, 122

newalias command, 122

components, 121

configure system to receive, 126

configuring system to send, 125

delivery agents, 121

sendmail command, 121

uux command, 121

MIME applications, 125

audio, 125

graphics, 125

video, 125

queue, 122

Reply-To\!: header, 125

SENDMAIL\_FREEZE environment variable, 125, 126, 127, 130

SENDMAIL\_SERVER environment variable, 125, 126, 127, 130

SENDMAIL\_SERVER\_NAME environment variable, 125, 126, 127, 130

TCP/IP addresses, 121

topography

central hub, 123, 126

choosing, 124

fully distributed, 124, 130

gateway hub, 123, 127

network, 123

standalone system, 130

user agents, 121

elm command, 121

mail command, 121

mailx command, 121

MIME applications, 121

UUCP addresses, 121

### mail files

\$HOME/.forward, 54

### maintenance

blocking nonsuperusers, 36

modify a user

HP SMH, text-based, 57  
usermod command, 60

mounting  
  file systems, 70

multisystem access  
  user ID, 69

multiuser mode, 38

**N**

Network File System (*see* NFS)

network gateway  
  setting with set\_parms command, 42

network parameters  
  setting, 41

networking, 75

NFS  
  /etc/dfs/dfstab file, 75  
  /usr/sbin/share command, 75  
  client  
    enable or disable, 76  
    enable or disable with HP SMH, text-based, 76  
  clients, 75  
  configuring, 75  
  exporting vs sharing, 75  
  server  
    enable or disable, 76  
    enable or disable with HP SMH, text-based, 76  
    share HP-UX directory, 77

servers, 75  
sharing, 75

NFS client  
  /etc/fstab file, 82, 84, 85  
  /etc/hosts file, 84, 85  
  /etc/rc.config.d/nfsconf file, 84  
  /etc/resolv.conf file, 84  
  /usr/sbin/mount command, 82  
  automount shared file system  
    with commands, 82  
    with HP SMH, text-based, 82  
  choosing automount mount, 78, 80  
  choosing standard mount, 78, 80  
  directory naming, 80  
  mount shared directory, 78  
  NFS\_CLIENT environment variable, 84  
  recovery after power failure, 86  
  required entries, 85  
  shared directory is renamed, 87  
  standard-mount shared file system  
    with HP SMH, text-based, 80  
  troubleshooting, 84

NFS server  
  /etc/dfs/dfstab file, 84, 85  
  /etc/rc.config.d/nfsconf file, 84  
  /usr/sbin/shareall command, 85  
  /usr/sbin/unshare command, 85  
  NFS\_CORE environment variable, 84  
  NFS\_SERVER environment variable, 84

recovery after power failure, 86  
renaming shared directory, 87

share directory  
  with commands, 77  
  with HP SMH, text-based, 77

troubleshooting, 84

NFS\_CLIENT environment variable, 84  
NFS\_CORE environment variable, 84  
NFS\_SERVER environment variable, 84

NIS server  
  setting with set\_parms command, 42

NOLOGIN variable, 36

**O**

operating system version, 13, 14

ownership  
  file, 68

**P**

passwd  
  comments field, 51  
  creating entry, 53  
  gecos field, 51  
  home directory field, 50  
  home phone field, 51  
  location field, 51  
  login name field, 50  
  login shell field, 51  
  password aging field, 52  
  password field, 53  
  phone field, 51  
  primary group field, 50  
  real name field, 51  
  user ID field, 50

password aging  
  passwd field, 52  
  shadow field, 52

password entry  
  creating, 53

passwords  
  aging, 52  
  passwd field, 53  
  restrictions, 52  
  setting, 53  
  shadow field, 53

PCI error recovery, 133  
  documentation, 135  
  HP Serviceguard, 133  
  pci\_eh\_enable kernel tunable, 134  
  pci\_error\_tolerance\_time kernel tunable, 134

pci\_eh\_enable kernel tunables, 134  
pci\_error\_tolerance\_time kernel tunable, 134

peripherals, 133  
  PCI OL\*, 133

permissions  
  file access, 68

phone  
    passwd field, 51  
plotters (*see* printers)  
power failure recovery, 86  
preloaded system  
    starting, 35  
primary group  
    passwd field, 50  
printers  
    adding class to spooler, 117  
    adding local printer to spooler, 112  
        with commands, 112  
        with HP SMH, text-based, 112  
    adding local printer to system, 112  
    adding network printer to spooler, 116  
        with commands, 116  
        with HP SMH, web-based, 116  
    adding printer class to spooler  
        with commands, 117  
    adding remote printer to spooler, 113  
        with commands, 114  
        with HP SMH, text-based, 113  
    default destination, 114  
LDAP, 111  
removing printer class, 120  
    with commands, 120  
removing printer from class, 119  
    with commands, 119  
removing printer from spooler, 118  
    with commands, 118  
    with HP SMH, web-based, 118  
summary, 111  
privileges  
    groups, 67  
publication history, 17

## Q

quit character, 37

## R

real name  
    passwd field, 51  
recovery after power failure, 86  
release information, 12  
release name, 13  
remote access, 69, 107, 110  
    sharing directories, 69, 70  
    Telnet, 107  
        from HP-UX to Windows, 108  
        from Windows to HP-UX, 109  
remote system access, 72  
restricted access  
    HP SMH, text-based  
        nonsuperusers, 31  
revision history, 17  
root

configuring, 47  
home directory, 47  
    permissions, 47  
passwd file, 47  
password  
    setting with set\_parms command, 42  
private files, 47  
run levels, 38

## S

security attributes  
    system default, 63  
    users, 65  
SENDMAIL\_FREEZE environment variable, 125, 126, 127, 130  
SENDMAIL\_SERVER environment variable, 125, 126, 127, 130  
SENDMAIL\_SERVER\_NAME environment variable, 125, 126, 127, 130  
servers  
    X Window System, 73  
shadow  
    login name field, 50  
    password aging field, 52  
    password field, 53  
sharing  
    file systems, 70  
shells  
    conventions, 16  
    login, 51  
    POSIX, 16  
single-user mode, 38  
skeleton directory, 54  
    changing with useradd command, 54  
    copying, 53  
SMB  
    sharing directories between HP-UX and Windows, 107  
SMH (*see* HP SMH)  
special characters  
    login, 37  
spooler  
    adding local printer, 112  
        with commands, 112  
        with HP SMH, text-based, 112  
    adding network printer, 116  
        with commands, 116  
        with HP SMH, web-based, 116  
    adding printer class, 117  
        with commands, 117  
    adding remote printer, 113  
        with commands, 114  
        with HP SMH, text-based, 113  
    removing printer, 118  
        with commands, 118  
        with HP SMH, web-based, 118  
    removing printer class, 120  
        with commands, 120  
    removing printer from class, 119  
        with commands, 119

starting, 111  
    commands, 112  
    HP SMH, text-based, 112  
stopping, 111  
    commands, 112  
    HP SMH, text-based, 112  
start character, 37  
startup files  
    \$HOME/.kshrc, 54  
    /etc/skel/.cshrc, 54  
    /etc/skel/.exrc, 54  
    /etc/skel/.login, 54  
    /etc/skel/.profile, 54  
stop character, 37  
subnetwork mask  
    setting with set\_parms command, 42  
superuser  
    shell, 16  
syntax  
    commands, 16  
    functions, 16  
system clock  
    problems, 40  
    setting, 40  
system default security attributes, 63  
System Management Homepage (*see* HP SMH)  
system parameters, 35  
    setting, 41  
system-wide login environment, 42  
systems  
    remote access, 72  
Systems Insight Manager (*see* HP SIM)

## T

template  
    HP SMH, text-based  
        create user, 55  
terminal emulation, 107, 110  
terminal special characters, 37  
terminals  
    configuring non-HP, 143  
    garbled display, 151  
    troubleshooting, 145  
    unresponsive, 145  
time  
    problems, 40  
    set\_parms command, 40  
    setting, 40  
    setting with set\_parms command, 42  
time zone  
    set\_parms command, 40  
    setting, 40  
    setting with set\_parms command, 42  
troubleshooting  
    NFS server and client, 84  
typographic conventions, 15

## U

unset shell command, 30  
user IDs  
    multiple systems, 69  
    case study, 70  
    passwd field, 50  
user login environment, 42  
user template  
    HP SMH, text-based  
        create, 55  
useradd command  
    add a user, 60  
    changing skeleton directory, 54  
userdel command  
    delete a user, 60  
usermod command  
    modify a user, 60  
users  
    account  
        aging, 51  
    add  
        command, 60  
        commands, 61  
        HP SMH, text-based, 50  
        useradd command, 60  
    comments, 51  
    delete  
        HP SMH, text-based, 58  
        userdel command, 60  
    home directory, 50  
    home phone, 51  
    location, 51  
    login name, 50  
    login shell, 51  
    modify  
        HP SMH, text-based, 57  
        usermod command, 60  
    password  
        aging, 52  
        restrictions, 52  
    phone, 51  
    primary group, 50  
    real name, 51  
    security attributes, 65  
    user ID, 50

## V

variables, 11  
    (*see also* environment variables)  
NOLOGIN, 36

## W

Windows  
    adding Windows system to HP-UX workgroup, 106  
    hardware connections, 106  
rlogin, 110

sharing directories between HP-UX and Windows, 106

sharing directories with CIFS (SMB), 107

Telnet

    from HP-UX to Windows, 108

    from Windows to HP-UX, 109

## X

X Window System, 32

    HP-UX client system, 32

    non-HP-UX client system, 33

server access, 73